



## **Cisco IOS XR Routing Configuration Guide**

Cisco IOS XR Software Release 3.4

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

Customer Order Number:  
Text Part Number: OL-14356-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

*Cisco IOS XR Routing Configuration Guide*  
© 2007 Cisco Systems, Inc. All rights reserved.



# CONTENTS

## **Preface RC-xiii**

- Changes to This Document **RC-xiii**
- Obtaining Documentation and Submitting a Service Request **RC-xiii**

## **Implementing BGP on Cisco IOS XR Software RC-1**

- Contents **RC-2**
- Prerequisites for Implementing BGP on Cisco IOS XR Software **RC-2**
- Information About Implementing BGP on Cisco IOS XR Software **RC-2**
  - BGP Functional Overview **RC-3**
  - BGP Router Identifier **RC-3**
  - BGP Default Limits **RC-4**
  - BGP Next Hop Tracking **RC-5**
  - Autonomous System Number Formats in BGP **RC-7**
  - BGP Configuration **RC-7**
  - No Default Address Family **RC-19**
  - Routing Policy Enforcement **RC-20**
  - Table Policy **RC-21**
  - Update Groups **RC-22**
  - BGP Cost Community **RC-22**
  - BGP Best Path Algorithm **RC-27**
  - Administrative Distance **RC-30**
  - Multiprotocol BGP **RC-32**
  - Route Dampening **RC-33**
  - BGP Routing Domain Confederation **RC-34**
  - BGP Route Reflectors **RC-34**
  - Default Address Family for show Commands **RC-37**
  - Distributed BGP **RC-37**
  - MPLS VPN Carrier Supporting Carrier **RC-38**
  - BGP Keychains **RC-39**
  - IPv6/IPv6 VPN Provider Edge Transport over MPLS **RC-39**
  - VPNv4/VPNv6 over the IP Core Using L2TPv3 Tunnels **RC-40**
- How to Implement BGP on Cisco IOS XR Software **RC-42**
  - Enabling BGP Routing **RC-44**
  - Configuring a Routing Domain Confederation for BGP **RC-47**

Resetting an eBGP Session Immediately Upon Link Failure	RC-49
Logging Neighbor Changes	RC-49
Adjusting BGP Timers	RC-50
Changing the BGP Default Local Preference Value	RC-51
Configuring the MED Metric for BGP	RC-53
Configuring BGP Weights	RC-54
Tuning the BGP Best-Path Calculation	RC-56
Indicating BGP Back-door Routes	RC-58
Configuring Aggregate Addresses	RC-60
Redistributing iBGP Routes into IGP	RC-61
Redistributing Prefixes into Multiprotocol BGP	RC-63
Configuring BGP Route Dampening	RC-66
Applying Policy When Updating the Routing Table	RC-71
Setting BGP Administrative Distance	RC-73
Configuring a BGP Neighbor Group and Neighbors	RC-74
Configuring a Route Reflector for BGP	RC-77
Configuring BGP Route Filtering by Route Policy	RC-79
Configuring BGP Next-hop Trigger Delay	RC-81
Disabling Next-hop Processing on BGP Updates	RC-83
Configuring BGP Community and Extended-Community Advertisements	RC-84
Configuring the BGP Cost Community	RC-86
Configuring Software to Store Updates from a Neighbor	RC-91
Configuring Distributed BGP	RC-93
Configuring a VPN Routing and Forwarding Instance in BGP	RC-96
Configuring Keychains for BGP	RC-114
Configuring an MDT Address Family Session in BGP	RC-115
Disabling a BGP Neighbor	RC-118
Resetting Neighbors Using BGP Inbound Soft Reset	RC-120
Resetting Neighbors Using BGP Outbound Soft Reset	RC-120
Resetting Neighbors Using BGP Hard Reset	RC-121
Clearing Caches, Tables, and Databases	RC-122
Displaying System and Network Statistics	RC-123
Displaying BGP Process Information	RC-125
Monitoring BGP Update Groups	RC-126
Configuration Examples for Implementing BGP on Cisco IOS XR Software	RC-127
Enabling BGP: Example	RC-127
Displaying BGP Update Groups: Example	RC-129
BGP Neighbor Configuration: Example	RC-129
BGP Confederation: Example	RC-130
BGP Route Reflector: Example	RC-131

BGP MDT Address Family Configuration: Example **RC-132**

Where to Go Next **RC-132**

Additional References **RC-133**

Related Documents **RC-133**

Standards **RC-133**

MIBs **RC-134**

RFCs **RC-134**

Technical Assistance **RC-134**

## **Implementing EIGRP on Cisco IOS XR Software** **RC-135**

Contents **RC-135**

Prerequisites for Implementing EIGRP on Cisco IOS XR Software **RC-136**

Restrictions for Implementing EIGRP on Cisco IOS XR Software **RC-136**

Information About Implementing EIGRP on Cisco IOS XR Software **RC-136**

EIGRP Functional Overview **RC-137**

EIGRP Features **RC-137**

EIGRP Components **RC-138**

EIGRP Configuration Grouping **RC-139**

EIGRP Configuration Modes **RC-139**

EIGRP Interfaces **RC-140**

Redistribution for an EIGRP Process **RC-140**

Metric Weights for EIGRP Routing **RC-141**

Percentage of Link Bandwidth Used for EIGRP Packets **RC-142**

Floating Summary Routes for an EIGRP Process **RC-142**

Split Horizon for an EIGRP Process **RC-144**

Adjustment of Hello Interval and Hold Time for an EIGRP Process **RC-145**

Stub Routing for an EIGRP Process **RC-145**

Route Policy Options for an EIGRP Process **RC-146**

EIGRP Layer 3 VPN PE-CE Site-of-Origin **RC-147**

IPv6 and IPv6 VPN Provider Edge Support over MPLS and IP **RC-148**

How to Implement EIGRP on Cisco IOS XR Software **RC-148**

Enabling EIGRP Routing **RC-149**

Configuring Route Summarization for an EIGRP Process **RC-151**

Redistributing Routes for EIGRP **RC-153**

Creating a Route Policy and Attaching It to an EIGRP Process **RC-155**

Configuring Stub Routing for an EIGRP Process **RC-158**

Configuring EIGRP as a PE-CE Protocol **RC-159**

Redistributing BGP Routes into EIGRP **RC-161**

Monitoring EIGRP Routing **RC-163**

- Configuration Examples for Implementing EIGRP on Cisco IOS XR Software **RC-166**
  - Configuring a Basic EIGRP Configuration: Example **RC-166**
  - Configuring an EIGRP Stub Operation: Example **RC-167**
  - Configuring an EIGRP PE-CE Configuration with Prefix-Limits: Example **RC-167**
- Additional References **RC-168**
  - Related Documents **RC-168**
  - Standards **RC-168**
  - MIBs **RC-168**
  - RFCs **RC-168**
  - Technical Assistance **RC-169**

**Implementing OSPF on Cisco IOS XR Software **RC-171****

- Contents **RC-172**
- Prerequisites for Implementing OSPF on Cisco IOS XR Software **RC-172**
- Information About Implementing OSPF on Cisco IOS XR Software **RC-173**
  - OSPF Functional Overview **RC-173**
  - Key Features Supported in the Cisco IOS XR OSPF Implementation **RC-175**
  - Comparison of Cisco IOS XR OSPFv3 and OSPFv2 **RC-175**
  - OSPF Hierarchical CLI and CLI Inheritance **RC-176**
  - OSPF Routing Components **RC-176**
  - OSPF Process and Router ID **RC-179**
  - Supported OSPF Network Types **RC-179**
  - Route Authentication Methods for OSPF **RC-180**
  - Neighbors and Adjacency for OSPF **RC-181**
  - Designated Router (DR) for OSPF **RC-181**
  - Default Route for OSPF **RC-181**
  - Link-State Advertisement Types for OSPF Version 2 **RC-181**
  - Link-State Advertisement Types for OSPFv3 **RC-182**
  - Virtual Link and Transit Area for OSPF **RC-183**
  - OSPFv2 Shamlink Support for MPLS VPN **RC-184**
  - Route Redistribution for OSPF **RC-186**
  - OSPF Shortest Path First Throttling **RC-186**
  - Nonstop Forwarding for OSPF Version 2 **RC-187**
  - Graceful Restart for OSPFv3 **RC-188**
  - Warm Standby and Nonstop Routing for OSPF Version 2 **RC-190**
  - Multicast-Intact Support for OSPF **RC-190**
  - Load Balancing in OSPF Version 2 and OSPFv3 **RC-191**
  - Multi-Area Adjacency for OSPF Version 2 **RC-191**
  - Label Distribution Protocol IGP Auto-configuration for OSPF **RC-192**
  - OSPF Authentication Message Digest Management **RC-193**

GTSM TTL Security Mechanism for OSPF	<b>RC-193</b>
Path Computation Element for OSPFv2	<b>RC-193</b>
How to Implement OSPF on Cisco IOS XR Software	<b>RC-194</b>
Enabling OSPF	<b>RC-194</b>
Configuring Stub and Not-So-Stubby Area Types	<b>RC-197</b>
Configuring Neighbors for Nonbroadcast Networks	<b>RC-199</b>
Configuring Authentication at Different Hierarchical Levels for OSPF Version 2	<b>RC-204</b>
Controlling the Frequency That the Same LSA Is Originated or Accepted for OSPF	<b>RC-207</b>
Creating a Virtual Link with MD5 Authentication to Area 0 for OSPF	<b>RC-209</b>
Summarizing Subnetwork LSAs on an OSPF ABR	<b>RC-213</b>
Redistributing Routes from One IGP into OSPF	<b>RC-215</b>
Configuring OSPF Shortest Path First Throttling	<b>RC-219</b>
Configuring Nonstop Forwarding Specific to Cisco for OSPF Version 2	<b>RC-224</b>
Configuring OSPF Version 2 for MPLS Traffic Engineering	<b>RC-226</b>
Configuring OSPFv3 Graceful Restart	<b>RC-230</b>
Configuring an OSPFv2 Sham Link	<b>RC-233</b>
Enabling Nonstop Routing for OSPFv2	<b>RC-236</b>
Enabling Multicast-intact for OSPFv2	<b>RC-237</b>
Associating Interfaces to a VRF	<b>RC-238</b>
Configuring OSPF as a Provider Edge to Customer Edge (PE-CE) Protocol	<b>RC-240</b>
Creating Multiple OSPF Instances (OSPF Process and a VRF)	<b>RC-243</b>
Configuring Multi-area Adjacency	<b>RC-244</b>
Configuring Label Distribution Protocol IGP Auto-configuration for OSPF	<b>RC-246</b>
Configuring Authentication Message Digest Management for OSPF	<b>RC-247</b>
Configuring Generalized TTL Security Mechanism (GTSM) for OSPF	<b>RC-251</b>
Verifying OSPF Configuration and Operation	<b>RC-254</b>
Configuration Examples for Implementing OSPF on Cisco IOS XR Software	<b>RC-255</b>
Cisco IOS XR for OSPF Version 2 Configuration: Example	<b>RC-255</b>
CLI Inheritance and Precedence for OSPF Version 2: Example	<b>RC-257</b>
MPLS TE for OSPF Version 2: Example	<b>RC-258</b>
ABR with Summarization for OSPFv3: Example	<b>RC-258</b>
ABR Stub Area for OSPFv3: Example	<b>RC-258</b>
ABR Totally Stub Area for OSPFv3: Example	<b>RC-258</b>
Route Redistribution for OSPFv3: Example	<b>RC-259</b>
Virtual Link Configured Through Area 1 for OSPFv3: Example	<b>RC-259</b>
Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example	<b>RC-260</b>
VPN Backbone and Sham Link Configured for OSPF Version 2: Example	<b>RC-260</b>
Where to Go Next	<b>RC-262</b>
Additional References	<b>RC-262</b>

Related Documents **RC-262**  
 Standards **RC-262**  
 MIBs **RC-263**  
 RFCs **RC-263**  
 Technical Assistance **RC-264**

**Implementing IS-IS on Cisco IOS XR Software** **RC-265**

Contents **RC-266**  
 Prerequisites for Implementing IS-IS on Cisco IOS XR Software **RC-266**  
 Restrictions for Implementing IS-IS on Cisco IOS XR Software **RC-266**  
 Information About Implementing IS-IS on Cisco IOS XR Software **RC-266**  
   IS-IS Functional Overview **RC-267**  
   Key Features Supported in the Cisco IOS XR IS-IS Implementation **RC-267**  
   IS-IS Configuration Grouping **RC-268**  
   IS-IS Configuration Modes **RC-268**  
   IS-IS Interfaces **RC-269**  
   Multitopology Configuration **RC-269**  
   IPv6 Routing and Configuring IPv6 Addressing **RC-269**  
   Limit LSP Flooding **RC-269**  
   Maximum LSP Lifetime and Refresh Interval **RC-270**  
   Overload Bit Configuration During Multitopology Operation **RC-270**  
   Single-Topology IPv6 Support **RC-271**  
   Multitopology IPv6 Support **RC-271**  
   IS-IS Authentication **RC-271**  
   Nonstop Forwarding **RC-272**  
   Multi-Instance IS-IS **RC-272**  
   Multiprotocol Label Switching Traffic Engineering **RC-273**  
   Overload Bit on Router **RC-273**  
   Default Routes **RC-273**  
   Attached Bit on an IS-IS Instance **RC-273**  
   IS-IS Support for Route Tags **RC-274**  
   Multicast-Intact Feature **RC-274**  
   Multicast Topology Support Using IS-IS **RC-275**  
   MPLS Label Distribution Protocol IGP Synchronization **RC-275**  
   Label Distribution Protocol IGP Auto-configuration **RC-276**  
   MPLS TE Forwarding Adjacency **RC-276**  
   MPLS TE Interarea Tunnels **RC-276**  
   IP Fast Reroute **RC-276**  
 How to Implement IS-IS on Cisco IOS XR Software **RC-277**  
   Enabling IS-IS and Configuring Level 1 or Level 2 Routing **RC-277**



Configuring Single Topology for IS-IS	RC-280
Configuring Multitopology for IS-IS	RC-284
Controlling LSP Flooding for IS-IS	RC-285
Configuring Nonstop Forwarding for IS-IS	RC-289
Configuring Authentication for IS-IS	RC-291
Configuring Keychains for IS-IS	RC-293
Configuring MPLS Traffic Engineering for IS-IS	RC-295
Tuning Adjacencies for IS-IS	RC-298
Setting SPF Interval for a Single-Topology IPv4 and IPv6 Configuration	RC-301
Customizing Routes for IS-IS	RC-303
Configuring MPLS LDP IS-IS Synchronization	RC-306
Enabling Multicast-Intact	RC-308
Tagging IS-IS Interface Routes	RC-309
Setting the Priority for Adding Prefixes to the RIB	RC-311
Configuring IP Fast Reroute Loop-free Alternate	RC-313
Configuration Examples for Implementing IS-IS on Cisco IOS XR Software	RC-314
Configuring Single-Topology IS-IS for IPv6: Example	RC-315
Configuring Multitopology IS-IS for IPv6: Example	RC-315
Redistributing IS-IS Routes Between Multiple Instances: Example	RC-315
Tagging Routes: Example	RC-316
Where to Go Next	RC-316
Additional References	RC-317
Related Documents	RC-317
Standards	RC-317
MIBs	RC-318
RFCs	RC-318
Technical Assistance	RC-318

## **Implementing and Monitoring RIB on Cisco IOS XR Software** RC-319

Contents	RC-320
Prerequisites for Implementing RIB on Cisco IOS XR Software	RC-320
Information About RIB Configuration	RC-320
Overview of RIB	RC-320
RIB Data Structures in BGP and Other Protocols	RC-321
RIB Administrative Distance	RC-321
RIB Support for IPv4 and IPv6	RC-322
RIB Statistics	RC-322
IPv6 and IPv6 VPN Provider Edge Transport over MPLS	RC-322
IP Fast Reroute	RC-323

- RIB Quarantining **RC-323**
- How to Deploy and Monitor RIB **RC-324**
  - Verifying RIB Configuration Using the Routing Table **RC-324**
  - Verifying Networking and Routing Problems **RC-324**
  - Disabling RIB Next-hop Dampening **RC-326**
- Configuration Examples for RIB Monitoring **RC-327**
  - Output of show route Command: Example **RC-328**
  - Output of show route backup Command: Example **RC-328**
  - Output of show route best-local Command: Example **RC-328**
  - Output of show route connected Command: Example **RC-329**
  - Output of show route local Command: Example **RC-329**
  - Output of show route longer-prefixes Command: Example **RC-329**
  - Output of show route next-hop Command: Example **RC-329**
- Where to Go Next **RC-330**
- Additional References **RC-330**
  - Related Documents **RC-331**
  - Standards **RC-331**
  - MIBs **RC-332**
  - RFCs **RC-332**
  - Technical Assistance **RC-332**
- Implementing RIP on Cisco IOS XR Software** **RC-333**
  - Contents **RC-333**
  - Information About Implementing RIP on Cisco IOS XR Software **RC-334**
    - Prerequisites for Implementing RIP on Cisco IOS XR Software **RC-334**
    - RIP Functional Overview **RC-334**
    - Split Horizon for RIP **RC-335**
    - Route Timers for RIP **RC-335**
    - Route Redistribution for RIP **RC-336**
    - Default Administrative Distances for RIP **RC-336**
    - Routing Policy Options for RIP **RC-337**
  - How to Implement RIP on Cisco IOS XR Software **RC-337**
    - Enabling RIP **RC-338**
    - Customize RIP **RC-339**
    - Control Routing Information **RC-342**
    - Creating a Route Policy for RIP **RC-344**
  - Configuration Examples for Implementing RIP on Cisco IOS XR Software **RC-347**
    - Configuring a Basic RIP Configuration: Example **RC-347**
    - Configuring RIP on the Provider Edge: Example **RC-348**

Adjusting RIP Timers for each VRF Instance: Example	RC-348
Configuring Redistribution for RIP: Example	RC-349
Configuring Route Policies for RIP: Example	RC-349
Configuring Passive Interfaces and Explicit Neighbors for RIP: Example	RC-350
Controlling RIP Routes: Example	RC-350
Additional References	RC-350
Related Documents	RC-351
Standards	RC-351
MIBs	RC-351
RFCs	RC-351
Technical Assistance	RC-352
<b>Implementing Routing Policy on Cisco IOS XR Software</b>	<b>RC-353</b>
Contents	RC-354
Prerequisites for Implementing Routing Policy on Cisco IOS XR Software	RC-354
Information About Implementing Routing Policy on Cisco IOS XR Software	RC-354
Routing Policy Language	RC-354
Routing Policy Language Overview	RC-355
Routing Policy Configuration Basics	RC-363
Policy Definitions	RC-363
Parameterization	RC-364
Semantics of Policy Application	RC-365
Policy Statements	RC-370
Attach Points	RC-374
Attached Policy Modification	RC-406
Nonattached Policy Modification	RC-406
How to Implement Routing Policy on Cisco IOS XR Software	RC-408
Defining a Route Policy	RC-409
Attaching a Routing Policy to a BGP Neighbor	RC-410
Modifying a Routing Policy Using a Text Editor	RC-412
Configuration Examples for Implementing Routing Policy on Cisco IOS XR Software	RC-413
Routing Policy Definition: Example	RC-414
Simple Inbound Policy: Example	RC-414
Modular Inbound Policy: Example	RC-415
Translating Cisco IOS Route Maps to Cisco IOS XR Routing Policy Language: Example	RC-416
Additional References	RC-416
Related Documents	RC-416
Standards	RC-417
MIBs	RC-417

RFCs **RC-417**  
 Technical Assistance **RC-417**

**Implementing Static Routes on Cisco IOS XR Software** **RC-419**

Contents **RC-420**

Prerequisites for Implementing Static Routes on Cisco IOS XR Software **RC-420**

Information About Implementing Static Routes on Cisco IOS XR Software **RC-420**

- Static Route Functional Overview **RC-420**
- Default Administrative Distance **RC-421**
- Directly Connected Routes **RC-421**
- Recursive Static Routes **RC-422**
- Fully Specified Static Routes **RC-422**
- Floating Static Routes **RC-423**
- Default VRF **RC-423**
- IPv4 and IPv6 Static VRF Routes **RC-423**
- IPv6/IPv6 VPN Provider Edge Transport over MPLS **RC-423**

How to Implement Static Routes on Cisco IOS XR Software **RC-424**

- Configuring a Static Route **RC-424**
- Configuring a Floating Static Route **RC-425**
- Configuring Static Routes Between PE-CE Routers **RC-427**
- Changing the Maximum Number of Allowable Static Routes **RC-429**
- Associating a VRF with a Static Route **RC-430**

Configuration Examples **RC-432**

- Configuring Traffic Discard: Example **RC-432**
- Configuring a Fixed Default Route: Example **RC-432**
- Configuring a Floating Static Route: Example **RC-432**
- Configuring a Static Route Between PE-CE Routers: Example **RC-432**

Where to Go Next **RC-433**

Additional References **RC-433**

- Related Documents **RC-433**
- Standards **RC-433**
- MIBs **RC-434**
- RFCs **RC-434**
- Technical Assistance **RC-434**

**Index**



## Preface

---

The Cisco IOS XR Routing Configuration Guide provides information and procedures related to routing on Cisco IOS XR software.

The preface contains the following sections:

- [Changes to This Document](#)
- [Obtaining Documentation and Submitting a Service Request](#)

## Changes to This Document

[Table 1](#) lists the technical changes made to this document since it was first printed.

**Table 1** *Changes to This Document*

Revision	Date	Change Summary
OL-14356-01	December 2007	Initial release of this document.

## Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.





# Implementing BGP on Cisco IOS XR Software

Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) that allows you to create loop-free interdomain routing between autonomous systems. An autonomous system is a set of routers under a single technical administration. Routers in an autonomous system can use multiple Interior Gateway Protocols (IGP) to exchange routing information inside the autonomous system and an EGP to route packets outside the autonomous system.

This module provides the conceptual and configuration information for BGP on Cisco IOS XR software.



## Note

For more information about BGP on the Cisco IOS XR software and complete descriptions of the BGP commands listed in this module, you can see the [“Related Documents”](#) section of this module. To locate documentation for other commands that might appear while performing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing BGP on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router.
Release 3.3.0	VPN routing and forwarding (VRF) support was added, including information on VRF command modes and command syntax. BGP cost community information was added.
Release 3.4.0	The following features were supported: <ul style="list-style-type: none"><li>• Four-byte autonomous system (AS) number</li><li>• Carrier supporting carrier (CSC) for BGP was added. See <i>Cisco IOS XR Multiprotocol Label Switching Protocol Configuration Guide</i> for information</li><li>• Key chains</li></ul>

Release 3.5.0	The following features were supported: <ul style="list-style-type: none"> <li>• IPv6 Provider Edge and IPv6 VPN Provider Edge over Multiprotocol Label Switching</li> <li>• Neighbor-specific VRF IPv6 address family configurations</li> <li>• Address family group-specific VPNv6 configurations</li> <li>• VPN4/VPNv6 over IP Core using L2TPv3 Tunnels</li> <li>• Multicast Distribution Tree (MDT) Subaddress Family Identifier Information (SAFI) support for Multicast VPN (MVPN)</li> </ul>
Release 3.6.0	No modification.

## Contents

- [Prerequisites for Implementing BGP on Cisco IOS XR Software, page RC-2](#)
- [Information About Implementing BGP on Cisco IOS XR Software, page RC-2](#)
- [How to Implement BGP on Cisco IOS XR Software, page RC-42](#)
- [Configuration Examples for Implementing BGP on Cisco IOS XR Software, page RC-127](#)
- [Where to Go Next, page RC-132](#)
- [Additional References, page RC-133](#)

## Prerequisites for Implementing BGP on Cisco IOS XR Software

You must be in a user group associated with a task group that includes the proper task IDs for BGP commands. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of *Cisco IOS XR System Security Configuration Guide*.

## Information About Implementing BGP on Cisco IOS XR Software

To implement BGP, you need to understand the following concepts:

- [BGP Functional Overview, page RC-3](#)
- [BGP Router Identifier, page RC-3](#)
- [BGP Default Limits, page RC-4](#)
- [BGP Next Hop Tracking, page RC-5](#)
- [Autonomous System Number Formats in BGP, page RC-7](#)
- [BGP Configuration, page RC-7](#)
- [No Default Address Family, page RC-19](#)
- [Routing Policy Enforcement, page RC-20](#)
- [Table Policy, page RC-21](#)
- [Update Groups, page RC-22](#)
- [BGP Best Path Algorithm, page RC-27](#)



- [Administrative Distance, page RC-30](#)
- [Administrative Distance, page RC-30](#)
- [Multiprotocol BGP, page RC-32](#)
- [Route Dampening, page RC-33](#)
- [BGP Routing Domain Confederation, page RC-34](#)
- [BGP Route Reflectors, page RC-34](#)
- [Default Address Family for show Commands, page RC-37](#)
- [Distributed BGP, page RC-37](#)
- [MPLS VPN Carrier Supporting Carrier, page RC-38](#)
- [BGP Keychains, page RC-39](#)
- [IPv6/IPv6 VPN Provider Edge Transport over MPLS, page RC-39](#)
- [VPNv4/VPNv6 over the IP Core Using L2TPv3 Tunnels, page RC-40](#)

## BGP Functional Overview

BGP uses TCP as its transport protocol. Two BGP routers form a TCP connection between one another (peer routers) and exchange messages to open and confirm the connection parameters.

BGP routers exchange network reachability information. This information is mainly an indication of the full paths (BGP autonomous system numbers) that a route should take to reach the destination network. This information helps construct a graph that shows which autonomous systems are loop free and where routing policies can be applied to enforce restrictions on routing behavior.

Any two routers forming a TCP connection to exchange BGP routing information are called peers or neighbors. BGP peers initially exchange their full BGP routing tables. After this exchange, incremental updates are sent as the routing table changes. BGP keeps a version number of the BGP table, which is the same for all of its BGP peers. The version number changes whenever BGP updates the table due to routing information changes. Keepalive packets are sent to ensure that the connection is alive between the BGP peers and notification packets are sent in response to error or special conditions.



**Note**

---

For information on configuring BGP to distribute Multiprotocol Label Switching (MPLS) Layer 3 virtual private network (VPN) information, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

---



**Note**

---

For information on BGP support for Bidirectional Forwarding Detection (BFD), see *Cisco IOS XR Interface and Hardware Configuration Guide* and *Cisco IOS XR Interface and Hardware Command Reference*.

---

## BGP Router Identifier

For BGP sessions between neighbors to be established, BGP must be assigned a router ID. The router ID is sent to BGP peers in the OPEN message when a BGP session is established.

BGP attempts to obtain a router ID in the following ways (in order of preference):

- By means of the address configured using the **bgp router-id** command in router configuration mode.
- By using the highest IPv4 address on a loopback interface in the system if the router is booted with saved loopback address configuration.
- By using the primary IPv4 address of the first loopback address that gets configured if there are not any in the saved configuration.

If none of these methods for obtaining a router ID succeeds, BGP does not have a router ID and cannot establish any peering sessions with BGP neighbors. In such an instance, an error message is entered in the system log, and the **show bgp summary** command displays a router ID of 0.0.0.0.

After BGP has obtained a router ID, it continues to use it even if a better router ID becomes available. This usage avoids unnecessary flapping for all BGP sessions. However, if the router ID currently in use becomes invalid (because the interface goes down or its configuration is changed), BGP selects a new router ID (using the rules described) and all established peering sessions are reset.


**Note**

We strongly recommend that the **bgp router-id** command is configured to prevent unnecessary changes to the router ID (and consequent flapping of BGP sessions).

## BGP Default Limits

Cisco IOS XR BGP imposes maximum limits on the number of neighbors that can be configured on the router and on the maximum number of prefixes that are accepted from a peer for a given address family. This limitation safeguards the router from resource depletion caused by misconfiguration, either locally or on the remote neighbor. The following limits apply to BGP configurations:

- The default maximum number of peers that can be configured is 4000. The default can be changed using the **bgp maximum neighbor** command. The *limit* range is 1 to 15000. Any attempt to configure additional peers beyond the maximum limit or set the maximum limit to a number that is less than the number of peers currently configured will fail.
- To prevent a peer from flooding BGP with advertisements, a limit is placed on the number of prefixes that are accepted from a peer for each supported address family. The default limits can be overridden through configuration of the **maximum-prefix limit** command for the peer for the appropriate address family. The following default limits are used if the user does not configure the maximum number of prefixes for the address family:
  - 512K (524,288) prefixes for IPv4 unicast.
  - 128K (131,072) prefixes for IPv4 multicast.
  - 128K (131,072) prefixes for IPv6 unicast.
  - 128K (131,072) prefixes for IPv6 multicast
  - 512K (524,288) prefixes for VPNv4 unicast
  - 512K (524,288) prefixes for VPNv6 unicast

A cease notification message is sent to the neighbor and the peering with the neighbor is terminated when the number of prefixes received from the peer for a given address family exceeds the maximum limit (either set by default or configured by the user) for that address family.

It is possible that the maximum number of prefixes for a neighbor for a given address family has been configured after the peering with the neighbor has been established and a certain number of prefixes have already been received from the neighbor for that address family. A cease notification message is sent to

the neighbor and peering with the neighbor is terminated immediately after the configuration if the configured maximum number of prefixes is fewer than the number of prefixes that have already been received from the neighbor for the address family.

## BGP Next Hop Tracking

BGP receives notifications from the Routing Information Base (RIB) when next-hop information changes (event-driven notifications). BGP obtains next-hop information from the RIB to:

- Determine whether a next hop is reachable.
- Find the fully recursed IGP metric to the next hop (used in the best-path calculation).
- Validate the received next hops.
- Calculate the outgoing next hops.
- Verify the reachability and connectedness of neighbors.

BGP is notified when any of the following events occurs:

- Next hop becomes unreachable
- Next hop becomes reachable
- Fully recursed IGP metric to the next hop changes
- First hop IP address or first hop interface change
- Next hop becomes connected
- Next hop becomes unconnected
- Next hop becomes a local address
- Next hop becomes a nonlocal address



### Note

---

Reachability and recursed metric events trigger a best-path recalculation.

---

Event notifications from the RIB are classified as critical and noncritical. Notifications for critical and noncritical events are sent in separate batches. However, a noncritical event is sent along with the critical events if the noncritical event is pending and there is a request to read the critical events.

- Critical events are related to the reachability (reachable and unreachable), connectivity (connected and unconnected), and locality (local and nonlocal) of the next hops. Notifications for these events are not delayed.
- Noncritical events include only the IGP metric changes. These events are sent at an interval of 3 seconds. A metric change event is batched and sent 3 seconds after the last one was sent.

The next-hop trigger delay for critical and noncritical events can be configured to specify a minimum batching interval for critical and noncritical events using the **nexthop trigger-delay** command. The trigger delay is address family dependent.

The BGP next-hop tracking feature allows you to specify that BGP routes are resolved using only next hops whose routes have the following characteristics:

- To avoid the aggregate routes, the prefix length must be greater than a specified value.
- The source protocol must be from a selected list, ensuring that BGP routes are not used to resolve next hops that could lead to oscillation.

This route policy filtering is possible because RIB identifies the source protocol of route that resolved a next hop as well as the mask length associated with the route. The **nexthop route-policy** command is used to specify the route-policy.

For information on route policy filtering for next hops using the next-hop attach point, see the *Implementing Routing Policy Language on Cisco IOS XR Software* module of *Cisco IOS XR Routing Configuration Guide*.

## Scoped IPv4/VPNv4 Table Walk

To determine which address family to process, a next-hop notification is received by first dereferencing the gateway context associated with the next hop, then looking into the gateway context to determine which address families are using the gateway context. The IPv4 unicast and VPNv4 unicast address families share the same gateway context, because they are registered with the IPv4 unicast table in the RIB. As a result, both the global IPv4 unicast table and the VPNv4 table are processed when an IPv4 unicast next-hop notification is received from the RIB. A mask is maintained in the next hop, indicating whether the next hop belongs to IPv4 unicast or VPNv4 unicast, or both. This scoped table walk localizes the processing in the appropriate address family table.

## Reordered Address Family Processing

The Cisco IOS XR software walks address family tables based on the numeric value of the address family. When a next-hop notification batch is received, the order of address family processing is reordered to the following order:

- IPv4 tunnel
- VPNv4 unicast
- VPNv6 unicast
- IPv4 labeled unicast
- IPv4 unicast
- IPv4 MDT
- IPv4 multicast
- IPv6 unicast
- IPv6 multicast
- IPv6 labeled unicast

## New Thread for Next-Hop Processing

The critical-event thread in the spkr process handles only next-hop, Bidirectional Forwarding Detection (BFD), and fast-external-failover (FEF) notifications. This critical-event thread ensures that BGP convergence is not adversely impacted by other events that may take a significant amount of time.

## show, clear, and debug Commands

The **show bgp nexthops** command provides statistical information about next-hop notifications, the amount of time spent in processing those notifications, and details about each next hop registered with the RIB. The **clear bgp nexthop performance-statistics** command ensures that the cumulative statistics associated with the processing part of the next-hop **show** command can be cleared to help in monitoring.

The **clear bgp nexthop registration** command performs an asynchronous registration of the next hop with the RIB. See the *BGP Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Commands* for information on the next-hop **show** and **clear** commands.

The **debug bgp nexthop** command displays information on next-hop processing. The **out** keyword provides debug information only about BGP registration of next hops with RIB. The **in** keyword displays debug information about next-hop notifications received from RIB. The **out** keyword displays debug information about next-hop notifications sent to the RIB. See *BGP Debug Commands on Cisco IOS XR Software*.

## Autonomous System Number Formats in BGP

Autonomous system numbers (ASNs) are globally unique identifiers used to identify autonomous systems (ASs) and enable ASs to exchange exterior routing information between neighboring ASs. A unique ASN is allocated to each AS for use in BGP routing.

Currently, ASNs are encoded as 2-byte numbers in BGP. The 2-byte range is 1 to 65535. To prepare for the eventual exhaustion of 2-byte ASNs, BGP has the capability to support 4-byte ASNs. The 4-byte range is 1.0 to 65535.65535 and the format is *high-order 16-bit value in decimal . low-order 16-bit value in decimal*. The BGP 4-byte ASN capability is used to propagate 4-byte-based AS path information across BGP speakers that do not support 4-byte AS numbers. This capability allows for the gradual transition from 2-byte ASNs to 4-byte ASNs. See draft-ietf-idr-as4bytes-12.txt for information on increasing the size of an ASN from 2 bytes to 4 bytes.

## BGP Configuration

Cisco IOS XR BGP follows a neighbor-based configuration model that requires that all configurations for a particular neighbor be grouped in one place under the neighbor configuration. Peer groups are not supported for either sharing configuration between neighbors or for sharing update messages. The concept of peer group has been replaced by a set of configuration groups to be used as templates in BGP configuration and automatically generated update groups to share update messages between neighbors. BGP configurations are grouped into four major categories:

- [Router Configuration Mode, page RC-8](#)
- [Router Address Family Configuration Mode, page RC-8](#)
- [Neighbor Configuration Mode, page RC-8](#)
- [Neighbor Address Family Configuration Mode, page RC-8](#)
- [VRF Configuration Mode, page RC-8](#)
- [VRF Address Family Configuration Mode, page RC-8](#)
- [VRF Neighbor Configuration Mode, page RC-8](#)
- [VRF Neighbor Address Family Configuration Mode, page RC-9](#)

## Configuration Modes

The following sections show how to enter some of the configuration modes. From a mode, you can enter the ? command to display the commands available in that mode.

## Router Configuration Mode

The following example shows how to enter router configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router (config)# router bgp 140
RP/0/RP0/CPU0:router (config-bgp)#
```

## Router Address Family Configuration Mode

The following example shows how to enter router address family configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 112
RP/0/RP0/CPU0:router (config-bgp)# address-family ipv4 multicast
RP/0/RP0/CPU0:router (config-bgp-af)#
```

## Neighbor Configuration Mode

The following example shows how to enter neighbor configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 140
RP/0/RP0/CPU0:router (config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router (config-bgp-nbr)#
```

## Neighbor Address Family Configuration Mode

The following example shows how to enter neighbor address family configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 112
RP/0/RP0/CPU0:router (config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router (config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-nbr-af)#
```

## VRF Configuration Mode

The following example shows how to enter VPN routing and forwarding (VRF) configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 140
RP/0/RP0/CPU0:router (config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router (config-bgp-vrf)#
```

## VRF Address Family Configuration Mode

The following example shows how to enter VRF address family configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 112
RP/0/RP0/CPU0:router (config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router (config-bgp-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-vrf-af)#
```

## VRF Neighbor Configuration Mode

The following example shows how to enter VRF neighbor configuration mode:

```
RP/0/RP0/CPU0:router (config)# router bgp 140
RP/0/RP0/CPU0:router (config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router (config-bgp-vrf)# neighbor 11.0.1.2
RP/0/RP0/CPU0:router (config-bgp-vrf-nbr)#
```

## VRF Neighbor Address Family Configuration Mode

The following example shows how to enter VRF neighbor address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 112
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 11.0.1.2
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#
```

## VPNv4 Address Family Configuration Mode

The following example shows how to enter VPNv4 address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 152
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#
```

## VPNv6 Address Family Configuration Mode

The following example shows how to enter VPNv6 address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 150
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#
```

## Neighbor Submode

Cisco IOS XR BGP uses a neighbor submode to make it possible to enter configurations without having to prefix every configuration with the **neighbor** keyword and the neighbor address:

- Cisco IOS XR software has a submode available for neighbors in which it is not necessary for every command to have a “neighbor *x.x.x.x*” prefix:

In Cisco IOS XR software, the configuration is as follows:

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.23.1.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
```

- An address family configuration submode inside the neighbor configuration submode is available for entering address family-specific neighbor configurations. In Cisco IOS XR, the configuration is as follows:

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 2002::2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2023
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy one in
```

- You must enter neighbor-specific IPv4, IPv6, VPNv4, or VPNv6 commands in neighbor address-family configuration submode. In Cisco IOS XR software, the configuration is as follows:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# maximum-prefix 1000
```

- You must enter neighbor-specific IPv4 and IPv6 commands in VRF neighbor address-family configuration submode. In Cisco IOS XR software, the configuration is as follows:

```
RP/0/RP0/CPU0:router(config)# router bgp 110
```

```
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 11.0.1.2
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)# route-policy pass all in
```

## Configuration Templates

The **af-group**, **session-group**, and **neighbor-group** configuration commands provide template support for the neighbor configuration in Cisco IOS XR software:

The **af-group** command is used to group address family-specific neighbor commands within an IPv4, IPv6, VPNv4, or VPNv6 address family. Neighbors that have the same address family configuration are able to use the address family group (af-group) name for their address family-specific configuration. A neighbor inherits the configuration from an address family group by way of the **use** command. If a neighbor is configured to use an address family group, the neighbor (by default) inherits the entire configuration from the address family group. However, a neighbor does not inherit all of the configuration from the address family group if items are explicitly configured for the neighbor. The address family group configuration is entered under the BGP router configuration mode. The following example shows how to enter address family group configuration mode.

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group afmcast1 address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)#
```

The **session-group** command allows you to create a session group from which neighbors can inherit address family-independent configuration. A neighbor inherits the configuration from a session group by way of the **use** command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the entire configuration of the session group. A neighbor does not inherit all of the configuration from a session group if a configuration is done directly on that neighbor. The following example shows how to enter session group configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)#
```

The **neighbor-group** command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address family groups and can comprise the complete configuration for a neighbor. After a neighbor group is configured, a neighbor can inherit the configuration of the group using the **use** command. If a neighbor is configured to use a neighbor group, the neighbor inherits the entire BGP configuration of the neighbor group.

The following example shows how to enter neighbor group configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 123
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group nbrgroup1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)#
```

The following example shows how to enter neighbor group address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group nbrgroup1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)#
```

- However, a neighbor does not inherit all of the configuration from the neighbor group if items are explicitly configured for the neighbor. In addition, some part of the configuration of the neighbor group could be hidden if a session group or address family group was also being used.



Configuration grouping has the following effects in Cisco IOS XR software:

- Commands entered at the session group level define address family-independent commands (the same commands as in the neighbor submode).
- Commands entered at the address family group level define address family-dependent commands for a specified address family (the same commands as in the neighbor-address family configuration submode).
- Commands entered at the neighbor group level define address family-independent commands and address family-dependent commands for each address family (the same as all available **neighbor** commands), and define the **use** command for the address family group and session group commands.

## Template Inheritance Rules

In Cisco IOS XR software, BGP neighbors or groups inherit configuration from other configuration groups.

For address family-independent configurations:

- Neighbors can inherit from session groups and neighbor groups.
- Neighbor groups can inherit from session groups and other neighbor groups.
- Session groups can inherit from other session groups.
- If a neighbor uses a session group and a neighbor group, the configurations in the session group are preferred over the global address family configurations in the neighbor group.

For address family-dependent configurations:

- Address family groups can inherit from other address family groups.
- Neighbor groups can inherit from address family groups and other neighbor groups.
- Neighbors can inherit from address family groups and neighbor groups.

Configuration group inheritance rules are numbered in order of precedence as follows:

1. If the item is configured directly on the neighbor, that value is used. In the example that follows, the advertisement interval is configured both on the neighbor group and neighbor configuration and the advertisement interval being used is from the neighbor configuration:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbr)# advertisement-interval 20
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 20 seconds:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.1.1.1

BGP neighbor is 10.1.1.1, remote AS 1, local AS 140, external link
Remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 20 seconds
```

```

For Address Family: IPv4 Unicast
BGP neighbor version 0
Update group: 0.1
eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
Route refresh request: received 0, sent 0
0 accepted prefixes
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:00:14, due to BGP neighbor initialized
External BGP neighbor not directly connected.

```

2. Otherwise, if the neighbor uses a session group or address family group, the configuration value is obtained from the session group or address family group. If the address family group or session group has a parent and an item is configured on the parent, the parent configuration is used. If the item is not configured on the parent but is configured on the parent of the parent, the configuration of the parent of the parent is used, and so on. In the example that follows, the advertisement interval is configured on a neighbor group and a session group and the advertisement interval value being used is from the session group:

```

RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group AS_2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 20
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group AS_2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1

```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```

RP/0/RP0/CPU0:router# show bgp neighbors 192.168.0.1

BGP neighbor is 192.168.0.1, remote AS 1, local AS 140, external link
Remote router ID 0.0.0.0
BGP state = Idle
Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
Received 0 messages, 0 notifications, 0 in queue
Sent 0 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 0
Update group: 0.1
eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
Route refresh request: received 0, sent 0
0 accepted prefixes
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:03:23, due to BGP neighbor initialized
External BGP neighbor not directly connected.

```

3. Otherwise, if the neighbor uses a neighbor group and does not use a session group or address family group, the configuration value can be obtained from the neighbor group either directly or through inheritance. In the example that follows, the advertisement interval from the neighbor group is used because it is not configured directly on the neighbor and no session group is used:

```
RP/0/RP0/CPU0:router(config)# router bgp 150
RP/0/RP0/CPU0:router(config-bgp)# session-group AS_2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 20
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```
RP/0/RP0/CPU0:router# show bgp neighbors 192.168.1.1

BGP neighbor is 192.168.2.2, remote AS 1, local AS 140, external link
Remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  Inbound path policy configured
  Policy for incoming advertisements is POLICY_1
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:01:14, due to BGP neighbor initialized
External BGP neighbor not directly connected.
```

To illustrate the same rule, the following example shows how to set the advertisement interval to 15 (from the session group) and 25 (from the neighbor group). The advertisement interval set in the session group overrides the one set in the neighbor group. The inbound policy is set to POLICY\_1 from the neighbor group.

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group ADV
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group ADV_2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 25
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy POLICY_1 in
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.2.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ADV
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group TIMER
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```
RP/0/RP0/CPU0:router# show bgp neighbors 192.168.2.2

BGP neighbor is 192.168.2.2, remote AS 1, local AS 140, external link
Remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:02:03, due to BGP neighbor initialized
External BGP neighbor not directly connected.
```

4. Otherwise, the default value is used. In the example that follows, neighbor 10.0.101.5 has the minimum time between advertisement runs set to 30 seconds (default) because the neighbor is not configured to use the neighbor configuration or the neighbor group configuration:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group adv_15
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 10
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.5
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.10
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group adv_15
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 30 seconds:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.5

BGP neighbor is 10.0.101.5, remote AS 1, local AS 140, external link
Remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 30 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.2
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
```

```

Route refresh request: received 0, sent 0
0 accepted prefixes
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
Connections established 0; dropped 0
Last reset 00:00:25, due to BGP neighbor initialized
External BGP neighbor not directly connected.

```

The inheritance rules used when groups are inheriting configuration from other groups are the same as the rules given for neighbors inheriting from groups.

## Template Inheritance

You can use the following **show** commands described to monitor BGP inheritance information:

- [show bgp neighbors](#), page RC-15
- [show bgp af-group](#), page RC-16
- [show bgp session-group](#), page RC-17
- [show bgp neighbor-group](#), page RC-18

### show bgp neighbors

Use the **show bgp neighbors** command to display information about the BGP configuration for neighbors.

- Use the **configuration** keyword to display the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or address family groups used by this neighbor.
- Use the **inheritance** keyword to display the session groups, neighbor groups, and address family groups from which this neighbor is capable of inheriting configuration.

The **show bgp neighbors** command examples that follow are based on the sample configuration:

```

RP/0/RP0/CPU0:router(config)# router bgp 142
RP/0/RP0/CPU0:router(config-bgp) # af-group GROUP_3 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp) # next-hop-self
RP/0/RP0/CPU0:router(config-bgp-afgrp) # route-policy POLICY_1 in
RP/0/RP0/CPU0:router(config-bgp-afgrp) # exit
RP/0/RP0/CPU0:router(config-bgp) # session-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-sngrp) # advertisement-interval 15
RP/0/RP0/CPU0:router(config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router(config-bgp) # neighbor-group GROUP_1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp) # use session-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp) # ebgp-multihop 3
RP/0/RP0/CPU0:router(config-bgp-nbrgrp) # address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af) # weight 100
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af) # send-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af) # exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp) # address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af) # default-originate
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af) # exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp) # exit
RP/0/RP0/CPU0:router(config-bgp) # neighbor 192.168.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr) # remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr) # use neighbor-group GROUP_1
RP/0/RP0/CPU0:router(config-bgp-nbr) # address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af) # use af-group GROUP_3
RP/0/RP0/CPU0:router(config-bgp-nbr-af) # weight 200

```

The following example displays sample output from the **show bgp neighbors** command using the **inheritance** keyword. The example shows that the neighbor inherits session parameters from neighbor group GROUP\_1, which in turn inherits from session group GROUP\_2. The neighbor inherits IPv4 unicast parameters from address family group GROUP\_3 and IPv4 multicast parameters from neighbor group GROUP\_1:

```
RP/0/RP0/CPU0:router# show bgp neighbors 192.168.0.1 inheritance

Session:          n:GROUP_1 s:GROUP_2
IPv4 Unicast:     a:GROUP_3
IPv4 Multicast:  n:GROUP_1
```

The following example displays sample output from the **show bgp neighbors** command using the **configuration** keyword. The example shows from where each item of configuration was inherited, or if it was configured directly on the neighbor (indicated by [ ]). For example, the **ebgp-multihop 3** command was inherited from neighbor group GROUP\_1 and the **next-hop-self** command was inherited from the address family group GROUP\_3:

```
RP/0/RP0/CPU0:router# show bgp neighbors 192.168.0.1 configuration

neighbor 192.168.0.1
  remote-as 2 []
  advertisement-interval 15 [n:GROUP_1 s:GROUP_2]
  ebgp-multihop 3 [n:GROUP_1]
  address-family ipv4 unicast []
  next-hop-self [a:GROUP_3]
  route-policy POLICY_1 in [a:GROUP_3]
  weight 200 []
  address-family ipv4 multicast [n:GROUP_1]
  default-originate [n:GROUP_1]
```

## show bgp af-group

Use the **show bgp af-group** command to display address family groups:

- Use the **configuration** keyword to display the effective configuration for the address family group, including any settings that have been inherited from address family groups used by this address family group.
- Use the **inheritance** keyword to display the address family groups from which this address family group is capable of inheriting configuration.
- Use the **users** keyword to display the neighbors, neighbor groups, and address family groups that inherit configuration from this address family group.

The **show bgp af-group** sample commands that follow are based on this sample configuration:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group GROUP_3 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RP0/CPU0:router(config-bgp-afgrp)# route-policy POLICY_1 in
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# af-group GROUP_1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# use af-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-afgrp)# maximum-prefix 2500 75 warning-only
RP/0/RP0/CPU0:router(config-bgp-afgrp)# default-originate
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# af-group GROUP_2 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# use af-group GROUP_3
RP/0/RP0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-afgrp)# capability orf prefix both
```

The following example displays sample output from the **show bgp af-group** command using the **configuration** keyword. This example shows from where each configuration item was inherited. The **default-originate** command was configured directly on this address family group (indicated by [ ]). The **remove-private-as** command was inherited from address family group GROUP\_2, which in turn inherited from address family group GROUP\_3:

```
RP/0/RP0/CPU0:router# show bgp af-group GROUP_1 configuration

af-group GROUP_1 address-family ipv4 unicast
  capability orf prefix-list both           [a:GROUP_2]
  default-originate                         [ ]
  maximum-prefix 2500 75 warning-only      [ ]
  route-policy POLICY_1 in                  [a:GROUP_2 a:GROUP_3]
  remove-private-AS                         [a:GROUP_2 a:GROUP_3]
  send-community-ebgp                       [a:GROUP_2]
  send-extended-community-ebgp             [a:GROUP_2]
```

The following example displays sample output from the **show bgp af-group** command using the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp af-group GROUP_2 users

IPv4 Unicast: a:GROUP_1
```

The following example displays sample output from the **show bgp af-group** command using the **inheritance** keyword. This shows that the specified address family group GROUP\_1 directly uses the GROUP\_2 address family group, which in turn uses the GROUP\_3 address family group:

```
RP/0/RP0/CPU0:router# show bgp af-group GROUP_1 inheritance

IPv4 Unicast: a:GROUP_2 a:GROUP_3
```

## show bgp session-group

Use the **show bgp session-group** command to display session groups:

- Use the **configuration** keyword to display the effective configuration for the session group, including any settings that have been inherited from session groups used by this session group.
- Use the **inheritance** keyword to display the session groups from which this session group is capable of inheriting configuration.
- Use the **users** keyword to display the session groups, neighbor groups, and neighbors that inherit configuration from this session group.

The output from the **show bgp session-group** command is based on the following session group configuration:

```
RP/0/RP0/CPU0:router(config)# router bgp 113
RP/0/RP0/CPU0:router(config-bgp)# session-group GROUP_1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# use session-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# update-source Loopback 0
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# session-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# use session-group GROUP_3
RP/0/RP0/CPU0:router(config-bgp-sngrp)# ebgp-multihop 2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# session-group GROUP_3
RP/0/RP0/CPU0:router(config-bgp-sngrp)# dmz-link-bandwidth
```

The following is sample output from the **show bgp session-group** command with the **configuration** keyword in EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp session-group GROUP_1 configuration

session-group GROUP_1
  ebgp-multihop 2          [s:GROUP_2]
  update-source Loopback0 []
  dmz-link-bandwidth      [s:GROUP_2 s:GROUP_3]
```

The following is sample output from the **show bgp session-group** command with the **inheritance** keyword showing that the GROUP\_1 session group inherits session parameters from the GROUP\_3 and GROUP\_2 session groups:

```
RP/0/RP0/CPU0:router# show bgp session-group GROUP_1 inheritance

Session: s:GROUP_2 s:GROUP_3
```

The following is sample output from the **show bgp session-group** command with the **users** keyword showing that both the GROUP\_1 and GROUP\_2 session groups inherit session parameters from the GROUP\_3 session group:

```
RP/0/RP0/CPU0:router# show bgp session-group GROUP_3 users

Session: s:GROUP_1 s:GROUP_2
```

## show bgp neighbor-group

Use the **show bgp neighbor-group** command to display neighbor groups:

- Use the **configuration** keyword to display the effective configuration for the neighbor group, including any settings that have been inherited from neighbor groups used by this neighbor group.
- Use the **inheritance** keyword to display the address family groups, session groups, and neighbor groups from which this neighbor group is capable of inheriting configuration.
- Use the **users** keyword to display the neighbors and neighbor groups that inherit configuration from this neighbor group.

The examples are based on the following group configuration:

```
RP/0/RP0/CPU0:router (config)# router bgp 140
RP/0/RP0/CPU0:router (config-bgp)# af-group GROUP_3 address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-afgrp)# remove-private-as
RP/0/RP0/CPU0:router (config-bgp-afgrp)# soft-reconfiguration inbound
RP/0/RP0/CPU0:router (config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router (config-bgp)# af-group GROUP_2 address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-afgrp)# use af-group GROUP_3
RP/0/RP0/CPU0:router (config-bgp-afgrp)# send-community-ebgp
RP/0/RP0/CPU0:router (config-bgp-afgrp)# send-extended-community-ebgp
RP/0/RP0/CPU0:router (config-bgp-afgrp)# capability orf prefix both
RP/0/RP0/CPU0:router (config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router (config-bgp)# session-group GROUP_3
RP/0/RP0/CPU0:router (config-bgp-sngrp)# timers 30 90
RP/0/RP0/CPU0:router (config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router (config-bgp)# neighbor-group GROUP_1
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# remote-as 1982
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# use neighbor-group GROUP_2
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router (config-bgp)# neighbor-group GROUP_2
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# use session-group GROUP_3
RP/0/RP0/CPU0:router (config-bgp-nbrgrp)# address-family ipv4 unicast
```



```
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# use af-group GROUP_2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# weight 100
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword. The configuration setting source is shown to the right of each command. In the output shown previously, the remote autonomous system is configured directly on neighbor group GROUP\_1, and the send community setting is inherited from neighbor group GROUP\_2, which in turn inherits the setting from address family group GROUP\_3:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group GROUP_1 configuration

neighbor-group GROUP_1
  remote-as 1982                               []
  timers 30 90                                 [n:GROUP_2 s:GROUP_3]
  address-family ipv4 unicast                  []
  capability orf prefix-list both              [n:GROUP_2 a:GROUP_2]
  remove-private-AS                           [n:GROUP_2 a:GROUP_2 a:GROUP_3]
  send-community-ebgp                          [n:GROUP_2 a:GROUP_2]
  send-extended-community-ebgp                 [n:GROUP_2 a:GROUP_2]
  soft-reconfiguration inbound                 [n:GROUP_2 a:GROUP_2 a:GROUP_3]
  weight 100                                   [n:GROUP_2]
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group GROUP\_1 inherits session (address family-independent) configuration parameters from neighbor group GROUP\_2. Neighbor group GROUP\_2 inherits its session parameters from session group GROUP\_3. It also shows that the GROUP\_1 neighbor group inherits IPv4 unicast configuration parameters from the GROUP\_2 neighbor group, which in turn inherits them from the GROUP\_2 address family group, which itself inherits them from the GROUP\_3 address family group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group GROUP_1 inheritance

Session:      n:GROUP-2 s:GROUP_3
IPv4 Unicast: n:GROUP_2 a:GROUP_2 a:GROUP_3
```

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the GROUP\_1 neighbor group inherits session (address family-independent) configuration parameters from the GROUP\_2 neighbor group. The GROUP\_1 neighbor group also inherits IPv4 unicast configuration parameters from the GROUP\_2 neighbor group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group GROUP_2 users

Session:      n:GROUP_1
IPv4 Unicast: n:GROUP_1
```

## No Default Address Family

BGP does not support the concept of a default address family. An address family must be explicitly configured under the BGP router configuration for the address family to be activated in BGP. Similarly, an address family must be explicitly configured under a neighbor for the BGP session to be activated under that address family. It is not required to have any address family configured under the BGP router configuration level for a neighbor to be configured. However, it is a requirement to have an address family configured at the BGP router configuration level for the address family to be configured under a neighbor.

## Routing Policy Enforcement

External BGP (eBGP) neighbors must have an inbound and outbound policy configured. If no policy is configured, no routes are accepted from the neighbor, nor are any routes advertised to it. This added security measure ensures that routes cannot accidentally be accepted or advertised in the case of a configuration omission error.



### Note

This enforcement affects only eBGP neighbors (neighbors in a different autonomous system than this router). For internal BGP (iBGP) neighbors (neighbors in the same autonomous system), all routes are accepted or advertised if there is no policy.

In the following example, for an eBGP neighbor, if all routes should be accepted and advertised with no modifications, a simple pass-all policy is configured:

```
RP/0/RP0/CPU0:router(config)# route-policy pass-all
RP/0/RP0/CPU0:router(config-rpl)# pass
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# commit
```

Use the **route-policy (BGP)** command in the neighbor address-family configuration mode to apply the pass-all policy to a neighbor. The following example shows how to allow all IPv4 unicast routes to be received from neighbor 192.168.40.42 and advertise all IPv4 unicast routes back to it:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 21
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all in
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit
```

Use the **show bgp summary** command to display eBGP neighbors that do not have both an inbound and outbound policy for every active address family. In the following example, such eBGP neighbors are indicated in the output with an exclamation (!) mark:

```
RP/0/RP0/CPU0:router# show bgp all all summary

Address Family: IPv4 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 41
BGP scan interval 60 secs
BGP is operating in STANDALONE mode.

Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          41            41           41

Neighbor         Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.1       0     1    919     925      41      0    0  15:15:08  10
10.0.101.2       0     2     0       0        0      0    0  00:00:00  Idle

Address Family: IPv4 Multicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 1
```

```
BGP scan interval 60 secs
BGP is operating in STANDALONE mode.
```

```
Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          1             1            1
```

Some configured eBGP neighbors do not have both inbound and outbound policies configured for IPv4 Multicast address family. These neighbors will default to sending and/or receiving no routes and are marked with '!' in the output below. Use the 'show bgp neighbor <nbr\_address>' command for details.

```
Neighbor        Spk   AS MsgRcvd MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.2      0     2     0         0         0     0     0 00:00:00 Idle!
```

```
Address Family: IPv6 Unicast
=====
```

```
BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 2
BGP scan interval 60 secs
BGP is operating in STANDALONE mode.
```

```
Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          2             2            2
```

```
Neighbor        Spk   AS MsgRcvd MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
2222::2         0     2     920       918         2     0     0 15:15:11 1
2222::4         0     3         0         0         0     0     0 00:00:00 Idle!
```

```
Address Family: IPv6 Multicast
=====
```

```
BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 1
BGP scan interval 60 secs
BGP is operating in STANDALONE mode.
```

```
Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          1             1            1
```

Some configured eBGP neighbors do not have both inbound and outbound policies configured for IPv6 Multicast address family. These neighbors will default to sending and/or receiving no routes and are marked with '!' in the output below. Use the 'show bgp neighbor <nbr\_address>' command for details.

```
Neighbor        Spk   AS MsgRcvd MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
2222::2         0     2     920       918         0     0     0 15:15:11 0
2222::4         0     3         0         0         0     0     0 00:00:00 Idle!
```

## Table Policy

The table policy feature in BGP allows you to configure traffic index values on routes as they are installed in the global routing table. This feature is enabled using the **table-policy** command and supports the BGP policy accounting feature.

BGP policy accounting uses traffic indices that are set on BGP routes to track various counters. See *Implementing Routing Policy on Cisco IOS XR Software* for details on table policy use. See the *Cisco Express Forwarding Commands on Cisco IOS XR Software* module in *Cisco IOS XR IP Addresses and Services Command Reference* for details on BGP policy accounting.

Table policy also provides the ability to drop routes from the RIB based on match criteria. This feature can be useful in certain applications and should be used with caution as it can easily create a routing ‘black hole’ where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

## Update Groups

The BGP Update Groups feature contains an algorithm that dynamically calculates and optimizes update groups of neighbors that share outbound policies and can share the update messages. The BGP Update Groups feature separates update group replication from peer group configuration, improving convergence time and flexibility of neighbor configuration.

To use this feature, you must understand the following concepts:

- [BGP Update Generation and Update Groups, page RC-22](#)
- [BGP Update Group, page RC-22](#)

## BGP Update Generation and Update Groups

The BGP Update Groups feature separates BGP update generation from neighbor configuration. The BGP Update Groups feature introduces an algorithm that dynamically calculates BGP update group membership based on outbound routing policies. This feature does not require any configuration by the network operator. Update group-based message generation occurs automatically and independently.

## BGP Update Group

When a change to the configuration occurs, the router automatically recalculates update group memberships and applies the changes.

For the best optimization of BGP update group generation, we recommend that the network operator keeps outbound routing policy the same for neighbors that have similar outbound policies. This feature contains commands for monitoring BGP update groups. For more information about the commands, see the “[Monitoring BGP Update Groups](#)” section on page RC-126.

## BGP Cost Community

The BGP cost community is a nontransitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best-path selection process by assigning cost values to specific routes. The extended community format defines generic points of insertion (POI) that influence the best-path decision at different points in the best-path algorithm.

The cost community attribute is applied to internal routes by configuring the **set extcommunity cost** command in a route policy. See the *Routing Policy Language Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Command Reference* for information on the **set extcommunity cost** command. The cost community set clause is configured with a cost community ID number (0–255) and cost community number (0–4294967295). The cost community number determines the preference for

the path. The path with the lowest cost community number is preferred. Paths that are not specifically configured with the cost community number are assigned a default cost community number of 2147483647 (the midpoint between 0 and 4294967295) and evaluated by the best-path selection process accordingly. When two paths have been configured with the same cost community number, the path selection process prefers the path with the lowest cost community ID. The cost-extended community attribute is propagated to iBGP peers when extended community exchange is enabled.

The following commands can be used to apply the route policy that is configured with the cost community set clause:

**route-policy at:**

- **aggregate-address**
- **redistribute**
- **network**

## How BGP Cost Community Influences the Best Path Selection Process

The cost community attribute influences the BGP best-path selection process at the point of insertion (POI). By default, the POI follows the Interior Gateway Protocol (IGP) metric comparison. When BGP receives multiple paths to the same destination, it uses the best-path selection process to determine which path is the best path. BGP automatically makes the decision and installs the best path in the routing table. The POI allows you to assign a preference to a specific path when multiple equal cost paths are available. If the POI is not valid for local best-path selection, the cost community attribute is silently ignored.

Cost communities are sorted first by POI then by community ID. Multiple paths can be configured with the cost community attribute for the same POI. The path with the lowest cost community ID is considered first. In other words, all cost community paths for a specific POI are considered, starting with the one with the lowest cost community. Paths that do not contain the cost community cost (for the POI and community ID being evaluated) are assigned the default community cost value (2147483647). If the cost community values are equal, then cost community comparison proceeds to the next lowest community ID for this POI.

To select the path with the lower cost community, simultaneously walk through the cost communities of both paths. This is done by maintaining two pointers to the cost community chain, one for each path, and advancing both pointers to the next applicable cost community at each step of the walk for the given POI, in order of community ID, and stop when a best path is chosen or the comparison is a tie. At each step of the walk, the following checks are done:

```
If neither pointer refers to a cost community,  
    Declare a tie;  
Elseif a cost community is found for one path but not for the other,  
    Choose the path with cost community as best path;  
Elseif the Community ID from one path is less than the other,  
    Choose the path with the lesser Community ID as best path;  
Elseif the Cost from one path is less than the other,  
    Choose the path with the lesser Cost as best path;  
Else Continue.
```



### Note

Paths that are not configured with the cost community attribute are considered by the best-path selection process to have the default cost value (half of the maximum value [4294967295] or 2147483647).

Applying the cost community attribute at the POI allows you to assign a value to a path originated or learned by a peer in any part of the local autonomous system or confederation. The cost community can be used as a “tie breaker” during the best-path selection process. Multiple instances of the cost community can be configured for separate equal cost paths within the same autonomous system or

confederation. For example, a lower cost community value can be applied to a specific exit path in a network with multiple equal cost exit points, and the specific exit path is preferred by the BGP best-path selection process. See the scenario described in the [“Influencing Route Preference in a Multiexit IGP Network” section on page RC-25](#).

**Note**

The cost community comparison in BGP is enabled by default. Use the **bgp bestpath cost-community ignore** command to disable the comparison.

See the [“BGP Best Path Algorithm” section on page RC-27](#) for information on the BGP best-path selection process.

## Cost Community Support for Aggregate Routes and Multipaths

The BGP cost community feature supports aggregate routes and multipaths. The cost community attribute can be applied to either type of route. The cost community attribute is passed to the aggregate or multipath route from component routes that carry the cost community attribute. Only unique IDs are passed, and only the highest cost of any individual component route is applied to the aggregate for each ID. If multiple component routes contain the same ID, the highest configured cost is applied to the route. For example, the following two component routes are configured with the cost community attribute using an inbound route policy:

- 10.0.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=100
- 192.168.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=200

If these component routes are aggregated or configured as a multipath, the cost value 200 is advertised, because it has the highest cost.

If one or more component routes do not carry the cost community attribute or the component routes are configured with different IDs, then the default value (2147483647) is advertised for the aggregate or multipath route. For example, the following three component routes are configured with the cost community attribute using an inbound route policy. However, the component routes are configured with two different IDs.

- 10.0.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=100
- 172.16.0.1
  - POI=IGP
  - cost community ID=2
  - cost number=100

- 192.168.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=200

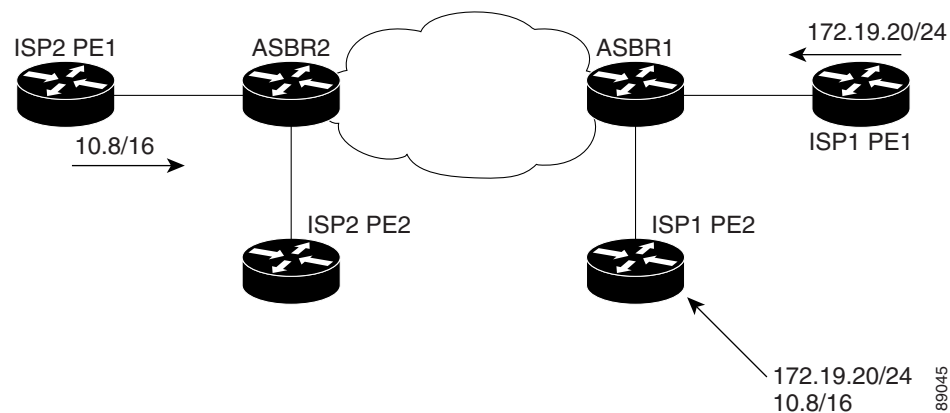
The single advertised path includes the aggregate cost communities as follows:

{POI=IGP, ID=1, Cost=2147483647} {POI=IGP, ID=2, Cost=2147483647}

## Influencing Route Preference in a Multiexit IGP Network

Figure 1 shows an IGP network with two autonomous system boundary routers (ASBRs) on the edge. Each ASBR has an equal cost path to network 10.8/16.

**Figure 1** Multiexit Point IGP Network



Both paths are considered to be equal by BGP. If multipath loadsharing is configured, both paths to the routing table are installed and are used to balance the load of traffic. If multipath load balancing is not configured, the BGP selects the path that was learned first as the best path and installs this path to the routing table. This behavior may not be desirable under some conditions. For example, the path is learned from ISP1 PE2 first, but the link between ISP1 PE2 and ASBR1 is a low-speed link.

The configuration of the cost community attribute can be used to influence the BGP best-path selection process by applying a lower-cost community value to the path learned by ASBR2. For example, the following configuration is applied to ASBR2:

```
RP/0/RP0/CPU0:router(config)# route-policy ISP2_PE1
RP/0/RP0/CPU0:router(config-rpl)# set extcommunity cost (1:1)
```

The preceding route policy applies a cost community number of 1 to the 10.8.0.0 route. By default, the path learned from ASBR1 is assigned a cost community number of 2147483647. Because the path learned from ASBR2 has a lower-cost community number, the path is preferred.

## BGP Cost Community Support for EIGRP MPLS VPN PE-CE with Back-door Links


Back-door links in an EIGRP MPLS VPN topology is preferred by BGP if the back-door link is learned first. (A back-door link, or route, is a connection that is configured outside of the VPN between a remote and main site; for example, a WAN leased line that connects a remote site to the corporate network.)

The “prebest path” point of insertion (POI) in the BGP cost community feature supports mixed EIGRP VPN network topologies that contain VPN and back-door links. This POI is applied automatically to EIGRP routes that are redistributed into BGP. The “prebest path” POI carries the EIGRP route type and metric. This POI influences the best-path calculation process by influencing BGP to consider the POI before any other comparison step. No configuration is required. This feature is enabled automatically for EIGRP VPN sites when Cisco IOS XR software is installed on a PE, CE, or back-door router.

For information about configuring EIGRP MPLS VPNs, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

Figure 2 shows how cost community can be used to support backdoor links in a network.

**Figure 2** Network Showing How Cost Community Can be Used to Support Backdoor Links



The following sequence of events happens in PE1:

1. PE1 learns IPv4 prefix 10.1.1.0/24 from CE1 through EIGRP running a virtual routing and forwarding (VRF) instance. EIGRP selects and installs the best path in the RIB. It also encodes the cost-extended community and adds the information to the RIB.
2. The route is redistributed into BGP (assuming that IGP-to-BGP redistribution is configured). BGP also receives the cost-extended community from the route through the redistribution process.
3. After BGP has determined the best path for the newly redistributed prefix, the path is advertised to PE peers (PE2).
4. PE2 receives the BGP VPNv4 prefix route\_distinguisher:10.1.1.0/24 along with the cost community. It is likely that CE2 advertises the same prefix (because of the back-door link between CE1 and CE2) to PE2 through EIGRP. PE2 BGP would have already learned the CE route through the redistribution process along with the cost community value.
5. PE2 has two paths within BGP: one with cost community cost1 through multipath BGP (PE1) and another with cost community cost2 through the EIGRP neighbor (CE2).
6. PE2 runs the enhanced BGP best-path calculation.
7. PE2 installs the best path in the RIB passing the appropriate cost community value.
8. PE2 RIB has two paths for 10.1.1.0/24: one with cost community cost2 added by EIGRP and another with the cost community cost1 added by BGP. Because both the route paths have cost community, RIB compares the costs first. The BGP path has the lower cost community, so it is selected and downloaded to the RIB.
9. PE2 RIB redistributes the BGP path into EIGRP with VRF. EIGRP runs a diffusing update algorithm (DUAL) because there are two paths, and selects the BGP-redistributed path.
10. PE2 EIGRP advertises the path to CE2 making the path the next hop for the prefix to send the traffic over the MPLS network.



## Adding Routes to the Routing Information Base

If a nonsourced path becomes the best path after the best-path calculation, BGP adds the route to the Routing Information Base (RIB) and passes the cost communities along with the other IGP extended communities.

When a route with paths is added to the RIB by a protocol, RIB checks the current best paths for the route and the added paths for cost extended communities. If cost-extended communities are found, the RIB compares the set of cost communities. If the comparison does not result in a tie, the appropriate best path is chosen. If the comparison results in a tie, the RIB proceeds with the remaining steps of the best-path algorithm. If a cost community is not present in either the current best paths or added paths, then the RIB continues with the remaining steps of the best-path algorithm. See the [“BGP Best Path Algorithm” section on page RC-27](#) for information on the BGP best-path algorithm.

## BGP Best Path Algorithm

BGP routers typically receive multiple paths to the same destination. The BGP best-path algorithm determines the best path to install in the IP routing table and to use for forwarding traffic. This section describes the Cisco IOS XR implementation of BGP best-path algorithm, as specified in Section 9.1 of the Internet Engineering Task Force (IETF) Network Working Group draft-ietf-idr-bgp4-24.txt document.

The BGP best-path algorithm implementation is in three parts:

- Part 1—Compares two paths to determine which is better.
- Part 2—Iterates over all paths and determines which order to compare the paths to select the overall best path.
- Part 3—Determines whether the old and new best paths differ enough so that the new best path should be used.

**Note**

The order of comparison determined by Part 2 is important because the comparison operation is not transitive; that is, if three paths, A, B, and C exist, such that when A and B are compared, A is better, and when B and C are compared, B is better, it is not necessarily the case that when A and C are compared, A is better. This nontransitivity arises because the multi exit discriminator (MED) is compared only among paths from the same neighboring autonomous system (AS) and not among all paths.

## Comparing Pairs of Paths

The following steps are completed to compare two paths and determine the better path:

1. If either path is invalid (for example, a path has the maximum possible MED value or it has an unreachable next hop), then the other path is chosen (provided that the path is valid).
2. If the paths have unequal pre-bestpath cost communities, the path with the lower pre-bestpath cost community is selected as the best path.

**Note**

See the [“BGP Cost Community” section on page RC-22](#) for details on how cost communities are compared.

3. If the paths have unequal weights, the path with the highest weight is chosen. Note: the weight is entirely local to the router, and can be set with the **weight** command or using a routing policy.
4. If the paths have unequal local preferences, the path with the higher local preference is chosen. Note: If a local preference attribute was received with the path or was set by a routing policy, then that value is used in this comparison. Otherwise, the default local preference value of 100 is used. The default value can be changed using the **bgp default local-preference** command.
5. If one of the paths is a redistributed path, which results from a **redistribute** or **network** command, then it is chosen. Otherwise, if one of the paths is a locally generated aggregate, which results from an **aggregate-address** command, it is chosen.

**Note**

Step 1 through Step 4 implement the “Degree of Preference” calculation from Section 9.1.1 of draft-ietf-idr-bgp4-24.txt.

6. If the paths have unequal AS path lengths, the path with the shorter AS path is chosen. This step is skipped if **bgp bestpath as-path ignore** command is configured. Note: when calculating the length of the AS path, confederation segments are ignored, and AS sets count as 1. (See Section 9.1.2.2a of draft-ietf-idr-bgp4-24.txt.)
7. If the paths have different origins, the path with the lower origin is selected. Interior Gateway Protocol (IGP) is considered lower than EGP, which is considered lower than INCOMPLETE. (See Section 9.1.2.2b of draft-ietf-idr-bgp4-24.txt.)
8. If appropriate, the MED of the paths is compared. If they are unequal, the path with the lower MED is chosen.

A number of configuration options exist that affect whether or not this step is performed. In general, the MED is compared if both paths were received from neighbors in the same AS; otherwise the MED comparison is skipped. However, this behavior is modified by certain configuration options, and there are also some corner cases to consider. (See Section 9.1.2.2c of draft-ietf-idr-bgp4-24.txt.)

If the **bgp bestpath med always** command is configured, then the MED comparison is always performed, regardless of neighbor AS in the paths. Otherwise, MED comparison depends on the AS paths of the two paths being compared, as follows:

- a. If a path has no AS path or the AS path starts with an AS\_SET, then the path is considered to be internal, and the MED is compared with other internal paths
- b. If the AS path starts with an AS\_SEQUENCE, then the neighbor AS is the first AS number in the sequence, and the MED is compared with other paths that have the same neighbor AS
- c. If the AS path contains only confederation segments or starts with confederation segments followed by an AS\_SET, then the MED is not compared with any other path unless the **bgp bestpath med confed** command is configured. In that case, the path is considered internal and the MED is compared with other internal paths.
- d. If the AS path starts with confederation segments followed by an AS\_SEQUENCE, then the neighbor AS is the first AS number in the AS\_SEQUENCE, and the MED is compared with other paths that have the same neighbor AS.

**Note**

If no MED attribute was received with the path, then the MED is considered to be 0 unless the **bgp bestpath med missing-as-worst** command is configured. In that case, if no MED attribute was received, the MED is considered to be the highest possible value.

9. If one path is received from an external peer and the other is received from an internal (or confederation) peer, the path from the external peer is chosen. (See Section 9.1.2.2d of draft-ietf-idr-bgp4-24.txt.)
10. If the paths have different IGP metrics to their next hops, the path with the lower IGP metric is chosen. (See Section 9.1.2.2e of draft-ietf-idr-bgp4-24.txt.)
11. If the paths have unequal IP cost communities, the path with the lower IP cost community is selected as the best path.

**Note**

See the “[BGP Cost Community](#)” section on page RC-22 for details on how cost communities are compared.

12. If all path parameters in Step 1 through Step 10 are the same, then the router IDs are compared. If the path was received with an originator attribute, then that is used as the router ID to compare; otherwise, the router ID of the neighbor from which the path was received is used. If the paths have different router IDs, the path with the lower router ID is chosen. Note: where the originator is used as the router ID, it is possible to have two paths with the same router ID. It is also possible to have two BGP sessions with the same peer router, and therefore receive two paths with the same router ID. (See Section 9.1.2.2f of draft-ietf-idr-bgp4-24.txt.)
13. If the paths have different cluster lengths, the path with the shorter cluster length is selected. If a path was not received with a cluster list attribute, it is considered to have a cluster length of 0.
14. Finally, the path received from the neighbor with the lower IP address is chosen. Locally generated paths (for example, redistributed paths) are considered to have a neighbor IP address of 0. (See Section 9.1.2.2g of draft-ietf-idr-bgp4-24.txt.)

## Order of Comparisons

The second part of the BGP best-path algorithm implementation determines the order in which the paths should be compared. The order of comparison is determined as follows:

1. The paths are partitioned into groups such that within each group the MED can be compared among all paths. The same rules as in the “[Comparing Pairs of Paths](#)” section on page RC-27 are used to determine whether MED can be compared between any two paths. Normally, this comparison results in one group for each neighbor AS. If the **bgp bestpath med always** command is configured, then there is just one group containing all the paths.
2. The best path in each group is determined. Determining the best path is achieved by iterating through all paths in the group and keeping track of the best one seen so far. Each path is compared with the best-so-far, and if it is better, it becomes the new best-so-far and is compared with the next path in the group.
3. A set of paths is formed containing the best path selected from each group in Step 2. The overall best path is selected from this set of paths, by iterating through them as in Step 2.

## Best Path Change Suppression

The third part of the implementation is to determine whether the best-path change can be suppressed or not—whether the new best path should be used, or continue using the existing best path. The existing best path can continue to be used if the new one is identical to the point at which the best-path selection algorithm becomes arbitrary (if the router-id is the same). Continuing to use the existing best path can avoid churn in the network.

**Note**

This suppression behavior does not comply with the IETF Networking Working Group draft-ietf-idr-bgp4-24.txt document, but is specified in the IETF Networking Working Group draft-ietf-idr-avoid-transition-00.txt document.

The suppression behavior can be turned off by configuring the **bgp bestpath compare-routerid** command. If this command is configured, the new best path is always preferred to the existing one.

Otherwise, the following steps are used to determine whether the best-path change can be suppressed:

1. If the existing best path is no longer valid, the change cannot be suppressed.
2. If either the existing or new best paths were received from internal (or confederation) peers or were locally generated (for example, by redistribution), then the change cannot be suppressed. That is, suppression is possible only if both paths were received from external peers.
3. If the paths were received from the same peer (the paths would have the same router-id), the change cannot be suppressed. The router ID is calculated using rules in the [“Comparing Pairs of Paths” section on page RC-27](#).
4. If the paths have different weights, local preferences, origins, or IGP metrics to their next hops, then the change cannot be suppressed. Note that all of these values are calculated using the rules in the [“Comparing Pairs of Paths” section on page RC-27](#).
5. If the paths have different-length AS paths and the **bgp bestpath as-path ignore** command is not configured, then the change cannot be suppressed. Again, the AS path length is calculated using the rules in the [“Comparing Pairs of Paths” section on page RC-27](#).
6. If the MED of the paths can be compared and the MEDs are different, then the change cannot be suppressed. The decision as to whether the MEDs can be compared is exactly the same as the rules in the [“Comparing Pairs of Paths” section on page RC-27](#), as is the calculation of the MED value.
7. If all path parameters in Step 1 through Step 6 do not apply, the change can be suppressed.

## Administrative Distance

An administrative distance is a rating of the trustworthiness of a routing information source. In general, the higher the value, the lower the trust rating. For information on specifying the administrative distance for BGP, see the *BGP Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Command Reference*.

Normally, a route can be learned through more than one protocol. Administrative distance is used to discriminate between routes learned from more than one protocol. The route with the lowest administrative distance is installed in the IP routing table. By default, BGP uses the administrative distances shown in [Table 1](#).

**Table 1** BGP Default Administrative Distances


Distance	Default Value	Function
External	20	Applied to routes learned from eBGP.
Internal	200	Applied to routes learned from iBGP.
Local	200	Applied to routes originated by the router.

**Note**

Distance does not influence the BGP path selection algorithm, but it does influence whether BGP-learned routes are installed in the IP routing table.

In most cases, when a route is learned through eBGP, it is installed in the IP routing table because of its distance (20). Sometimes, however, two ASs have an IGP-learned back-door route and an eBGP-learned route. Their policy might be to use the IGP-learned path as the preferred path and to use the eBGP-learned path when the IGP path is down. See [Figure 3](#).

**Figure 3**      **Back Door Example**



149547

In [Figure 3](#), Routers A and C and Routers B and C are running eBGP. Routers A and B are running an IGP (such as Routing Information Protocol [RIP], Interior Gateway Routing Protocol [IGRP], Enhanced IGRP, or Open Shortest Path First [OSPF]). The default distances for RIP, IGRP, Enhanced IGRP, and OSPF are 120, 100, 90, and 110, respectively. All these distances are higher than the default distance of eBGP, which is 20. Usually, the route with the lowest distance is preferred.

Router A receives updates about 160.10.0.0 from two routing protocols: eBGP and IGP. Because the default distance for eBGP is lower than the default distance of the IGP, Router A chooses the eBGP-learned route from Router C. If you want Router A to learn about 160.10.0.0 from Router B (IGP), establish a BGP back door. See [“Indicating BGP Back-door Routes”](#) section on page RC-58.

In the following example, a network back-door is configured:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# network 160.10.0.0/16 backdoor
```

Router A treats the eBGP-learned route as local and installs it in the IP routing table with a distance of 200. The network is also learned through Enhanced IGRP (with a distance of 90), so the Enhanced IGRP route is successfully installed in the IP routing table and is used to forward traffic. If the Enhanced IGRP-learned route goes down, the eBGP-learned route is installed in the IP routing table and is used to forward traffic.

Although BGP treats network 160.10.0.0 as a local entry, it does not advertise network 160.10.0.0 as it normally would advertise a local entry.




In [Figure 5](#), only unicast traffic can travel from Router A to the unicast routers to Router B and back. Multicast traffic could not flow on that path, so another routing table is required. Multicast traffic uses the path from Router A to the multicast routers to Router B and back.

[Figure 5](#) illustrates a multiprotocol BGP environment with a separate unicast route and multicast route from Router A to Router B. Multiprotocol BGP allows these routes to be incongruent. Both of the autonomous systems must be configured for internal multiprotocol BGP (IMBGP) in the figure.

A multicast routing protocol, such as PIM, uses the multicast BGP database to perform Reverse Path Forwarding (RPF) lookups for multicast-capable sources. Thus, packets can be sent and accepted on the multicast topology but not on the unicast topology.

**Figure 5** *Multicast BGP Environment*



11754

## Route Dampening

Route dampening is a BGP feature that minimizes the propagation of flapping routes across an internetwork. A route is considered to be flapping when it is repeatedly available, then unavailable, then available, then unavailable, and so on.

For example, consider a network with three BGP autonomous systems: autonomous system 1, autonomous system 2, and autonomous system 3. Suppose the route to network A in autonomous system 1 flaps (it becomes unavailable). Under circumstances without route dampening, the eBGP neighbor of autonomous system 1 to autonomous system 2 sends a withdraw message to autonomous system 2. The border router in autonomous system 2, in turn, propagates the withdrawal message to autonomous system 3. When the route to network A reappears, autonomous system 1 sends an advertisement message

to autonomous system 2, which sends it to autonomous system 3. If the route to network A repeatedly becomes unavailable, then available, many withdrawal and advertisement messages are sent. Route flapping is a problem in an internetwork connected to the Internet, because a route flap in the Internet backbone usually involves many routes.

## Minimizing Flapping

The route dampening feature minimizes the flapping problem as follows. Suppose again that the route to network A flaps. The router in autonomous system 2 (in which route dampening is enabled) assigns network A a penalty of 1000 and moves it to history state. The router in autonomous system 2 continues to advertise the status of the route to neighbors. The penalties are cumulative. When the route flaps so often that the penalty exceeds a configurable suppression limit, the router stops advertising the route to network A, regardless of how many times it flaps. Thus, the route is dampened.

The penalty placed on network A is decayed until the reuse limit is reached, upon which the route is once again advertised. At half of the reuse limit, the dampening information for the route to network A is removed.

**Note**

---

No penalty is applied to a BGP peer reset when route dampening is enabled, even though the reset withdraws the route.

---

## BGP Routing Domain Confederation

One way to reduce the iBGP mesh is to divide an autonomous system into multiple subautonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Although the peers in different autonomous systems have eBGP sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, MED, and local preference information is preserved. This feature allows you to retain a single IGP for all of the autonomous systems.

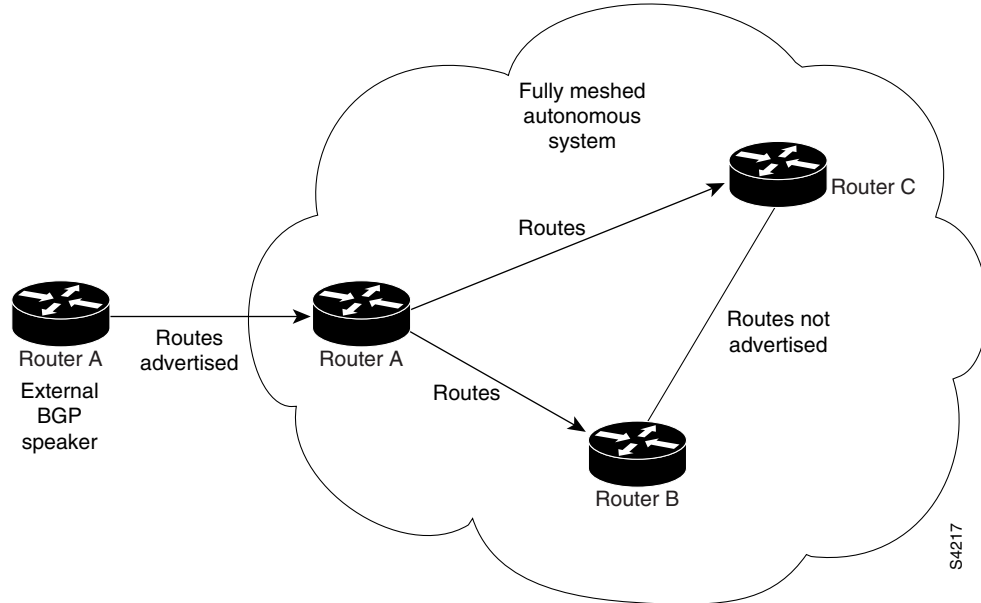
## BGP Route Reflectors

BGP requires that all iBGP speakers be fully meshed. However, this requirement does not scale well when there are many iBGP speakers. Instead of configuring a confederation, you can reduce the iBGP mesh by using a route reflector configuration.

[Figure 6](#) illustrates a simple iBGP configuration with three iBGP speakers (routers A, B, and C). Without route reflectors, when Router A receives a route from an external neighbor, it must advertise it to both routers B and C. Routers B and C do not readvertise the iBGP learned route to other iBGP speakers because the routers do not pass on routes learned from internal neighbors to other internal neighbors, thus preventing a routing information loop.

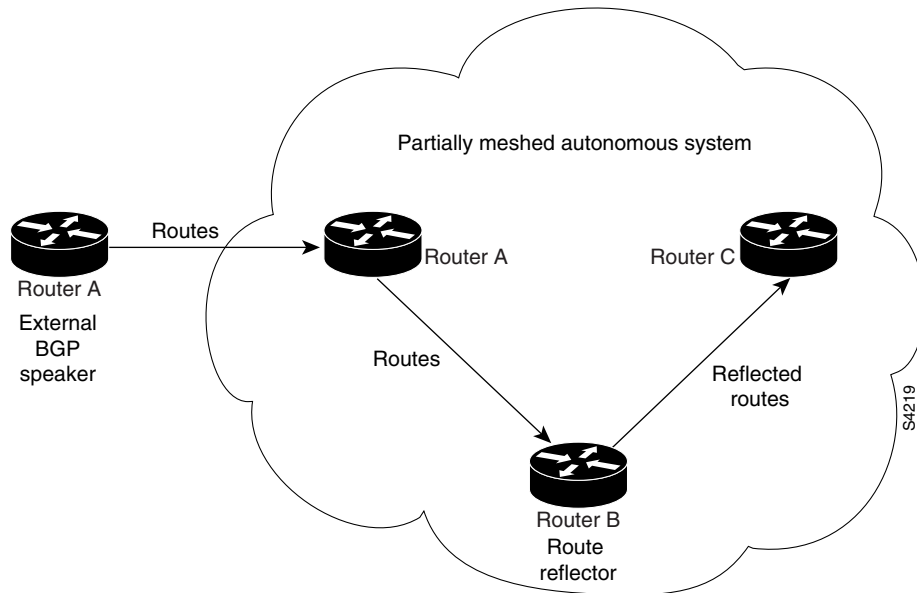


**Figure 6** Three Fully Meshed iBGP Speakers



With route reflectors, all iBGP speakers need not be fully meshed because there is a method to pass learned routes to neighbors. In this model, an iBGP peer is configured to be a route reflector responsible for passing iBGP learned routes to a set of iBGP neighbors. In Figure 7, Router B is configured as a route reflector. When the route reflector receives routes advertised from Router A, it advertises them to Router C, and vice versa. This scheme eliminates the need for the iBGP session between routers A and C.

**Figure 7** Simple BGP Model with a Route Reflector



The internal peers of the route reflector are divided into two groups: client peers and all other routers in the autonomous system (nonclient peers). A route reflector reflects routes between these two groups. The route reflector and its client peers form a *cluster*. The nonclient peers must be fully meshed with each other, but the client peers need not be fully meshed. The clients in the cluster do not communicate with iBGP speakers outside their cluster.

**Figure 8** *More Complex BGP Route Reflector Model*

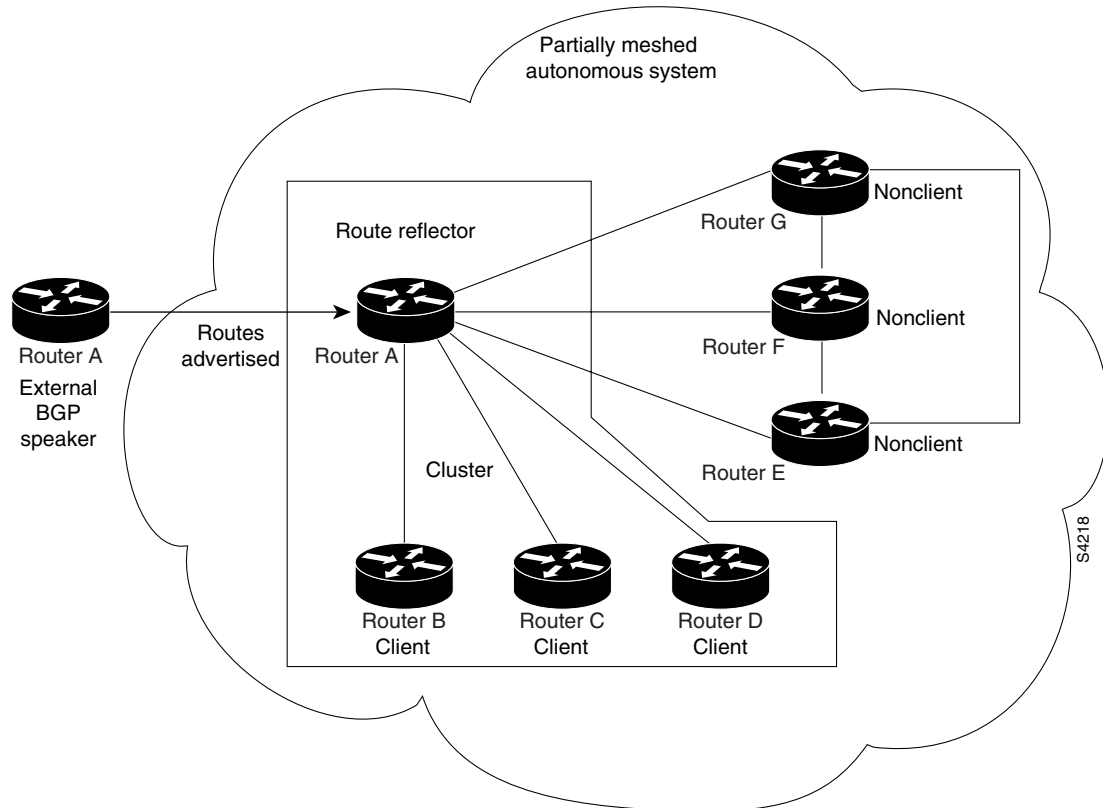


Figure 8 illustrates a more complex route reflector scheme. Router A is the route reflector in a cluster with routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

When the route reflector receives an advertised route, depending on the neighbor, it takes the following actions:

- A route from an external BGP speaker is advertised to all clients and nonclient peers.
- A route from a nonclient peer is advertised to all clients.
- A route from a client is advertised to all clients and nonclient peers. Hence, the clients need not be fully meshed.

Along with route reflector-aware BGP speakers, it is possible to have BGP speakers that do not understand the concept of route reflectors. They can be members of either client or nonclient groups, allowing an easy and gradual migration from the old BGP model to the route reflector model. Initially, you could create a single cluster with a route reflector and a few clients. All other iBGP speakers could be nonclient peers to the route reflector and then more clusters could be created gradually.

An autonomous system can have multiple route reflectors. A route reflector treats other route reflectors just like other iBGP speakers. A route reflector can be configured to have other route reflectors in a client group or nonclient group. In a simple configuration, the backbone could be divided into many clusters.

Each route reflector would be configured with other route reflectors as nonclient peers (thus, all route reflectors are fully meshed). The clients are configured to maintain iBGP sessions with only the route reflector in their cluster.

Usually, a cluster of clients has a single route reflector. In that case, the cluster is identified by the router ID of the route reflector. To increase redundancy and avoid a single point of failure, a cluster might have more than one route reflector. In this case, all route reflectors in the cluster must be configured with the cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. All route reflectors serving a cluster should be fully meshed and all of them should have identical sets of client and nonclient peers.

By default, the clients of a route reflector are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, the route reflector need not reflect routes to clients.

As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:

- Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector. The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come back to the originator, the information is ignored.
- Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

## Default Address Family for show Commands

Most of the **show** commands provide address family (AFI) and subaddress family (SAFI) arguments (see RFC 1700 and RFC 2858 for information on AFI and SAFI). The Cisco IOS XR software parser provides the ability to set the **afi** and **safi** so that it is not necessary to specify them while running a **show** command. The parser commands are:

- **set default-afi {ipv4 | ipv6 | all}**
- **set default-safi {unicast | multicast | all}**

The parser automatically sets the default **afi** value to **ipv4** and default **safi** value to **unicast**. It is necessary to use only the parser commands to change the default **afi** value from **ipv4** or default **safi** value from **unicast**. Any **afi** or **safi** keyword specified in a **show** command overrides the values set using the parser commands. Use the following command to check the currently set value of the **afi** and **safi**:

- **show default-afi-safi-vrf**

## Distributed BGP

Distributed BGP splits BGP functionality into three process types:

- BGP process manager—Responsible for verifying configuration changes and for calculating and publishing the distribution of neighbors among BGP speaker processes.

There is a single instance of this process.

- **bRIB process**—Responsible for performing the best-path calculation of routes (receives partial best paths from the speaker). The best route is installed into the bRIB and is advertised back to all speakers. See “[BGP Best Path Algorithm](#)” section on page RC-27 for information on best-path calculation. The bRIB process is also responsible for installing routes in the RIB, and for handling routes redistributed from the RIB. To accommodate route leaking from one RIB to another, bRIB may register for redistribution from multiple RIB routes into a single route in the bRIB process.

There is a single instance of this process for each address family.

- **BGP speaker process**—Responsible for handling all BGP connections to peers. The speaker stores received paths in the RIB and performs a partial best-path calculation, advertising the partial best paths to the bRIB (limited best-path calculation). Speakers perform a limited best-path calculation because to compare Multi Exit Discriminators (MEDs), paths need to be compared from the same AS but may not be received on the same speaker. Because BGP speakers do not have access to the entire BGP local RIB, BGP speakers can perform only a limited best-path calculation. (These are Step 1 through Step 7 in the “[BGP Best Path Algorithm](#)” section on page RC-27.) Only the best paths are advertised to the bRIB to reduce speaker/bRIB interprocess communications (IPC) and to reduce the number of paths to be processed in the bRIB. BGP speakers can only mark a path as active only after learning the result of the full best-path calculation from the bRIB. Neighbor import and export policies are imposed by the speaker.

If the **bgp bestpath med always** command is enabled, complete best-path calculation happens inside speaker process. When the **bgp bestpath med always** command is not enabled, speakers calculate partial best paths only (performs the best-path steps up to the MED comparison) and send them to bRIB. bRIB calculates the final best path (performs all the steps in the best-path calculation). When the **bgp bestpath med always** command is enabled, speakers can compare the MED across all ASs, allowing the speaker to calculate a single best path to send it to bRIB. bRIB is the ultimate process that calculates the final best path, but when the **bgp bestpath med always** command is enabled, the speakers send a single best path instead of potentially sending multiple partial best paths.

There are multiple instances of this process in which each instance is responsible for a subset of BGP peer connections.

Up to a total 15 speakers for all address families and one bRIB for each address family (IPv4, IPv6, and VPNv4) are supported.

Distributed BGP is used to reduce the impact that a fault in one address family has on another address family. For example, you can have one speaker with only IPv6 neighbors (peering to IPv6 addresses) and a separate speaker with only IPv4 neighbors (peering to IPv4 addresses), and yet another speaker with only VPNv4 provider edge (PE) or customer edge (CE) neighbors (peering to IPv4 addresses distinct from the non-VPN neighbors). In this scenario, there is no overlap in processes (bgp, bRIB, and RIB) between IPv4, IPv6, and VPNv4. Therefore, a bgp, bRIB, or RIB process crash affects only one address family. Distributed BGP also allows more CPU capacity for receiving, computing, and sending BGP routing updates. When in distributed BGP mode, you can control the number of distributed speakers that are enabled, as well as which neighbors are assigned to each speaker. If no distributed speakers are enabled, BGP operates in standalone mode. If at least one distributed speaker is enabled, BGP operates in distributed mode.

## MPLS VPN Carrier Supporting Carrier

Carrier supporting carrier (CSC) is a term used to describe a situation in which one service provider allows another service provider to use a segment of its backbone network. The service provider that provides the segment of the backbone network to the other provider is called the *backbone carrier*. The service provider that uses the segment of the backbone network is called the *customer carrier*.

A backbone carrier offers Border Gateway Protocol and Multiprotocol Label Switching (BGP/MPLS) VPN services. The customer carrier can be either:

- An Internet service provider (ISP) (By definition, an ISP does not provide VPN service.)
- A BGP/MPLS VPN service provider

You can configure a CSC network to enable BGP to transport routes and MPLS labels between the backbone carrier provider edge (PE) routers and the customer carrier customer edge (CE) routers using multiple paths. The benefits of using BGP to distribute IPv4 routes and MPLS label routes are:

- BGP takes the place of an Interior Gateway Protocol (IGP) and Label Distribution Protocol (LDP) in a VPN routing and forwarding (VRF) table. You can use BGP to distribute routes and MPLS labels. Using a single protocol instead of two simplifies the configuration and troubleshooting.
- BGP is the preferred routing protocol for connecting two ISPs, mainly because of its routing policies and ability to scale. ISPs commonly use BGP between two providers. This feature enables those ISPs to use BGP.

For detailed information on configuring MPLS VPN CSC with BGP, see the *Implementing MPLS Layer 3 VPNs on Cisco IOS XR Software* module of *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## BGP Keychains

BGP keychains enable keychain authentication between two BGP peers. The BGP endpoints must both comply with draft-bonica-tcp-auth-05.txt and a keychain on one endpoint and a password on the other endpoint does not work.

See *Cisco IOS XR System Security Guide* for information on keychain management.

BGP is able to use the keychain to implement hitless key rollover for authentication. The key rollover specification is time based, and in the event of clock skew between the peers, the rollover process is impacted. The configurable tolerance specification allows for the accept window to be extended (before and after) by that margin. This accept window facilitates a hitless key rollover for applications (for example, routing and management protocols).

The key rollover does not impact the BGP session, unless there is a keychain configuration mismatch at the endpoints resulting in no common keys for the session traffic (send or accept).

## IPv6/IPv6 VPN Provider Edge Transport over MPLS

IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) leverages the existing Multiprotocol Label Switching (MPLS) IPv4 core infrastructure for IPv6 transport. 6PE and 6VPE enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs). This feature relies on multiprotocol Border Gateway Protocol (BGP) extensions in the IPv4 network configuration on the provider edge (PE) router, to exchange IPv6 reachability information in addition to an MPLS label for each IPv6 address prefix to be advertised. Edge routers are configured to be dual stack running both IPv4 and IPv6, and use the IPv4-mapped IPv6 address for IPv6 prefix reachability exchange.



### Note

This feature is supported on Cisco XR 12000 Series Routers.

For detailed information on configuring 6PE and 6VPE over MPLS, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## IPv6 Provider Edge Multipath

Internal and external BGP multipath for IPv6 allows the IPv6 router to load balance between several paths (for example, same neighboring autonomous system [AS] or sub-AS, or the same metric) to reach its destination. The 6PE multipath feature uses multiprotocol internal BGP (MP-iBGP) to distribute IPv6 routes over the MPLS IPv4 core network and to attach an MPLS label to each route.

When MP-iBGP multipath is enabled on the 6PE router, all labeled paths are installed in the forwarding table with MPLS information (label stack) when MPLS information is available. This functionality enables 6PE to perform load balancing.

## VPNv4/VPNv6 over the IP Core Using L2TPv3 Tunnels

The Layer 2 Tunnel Protocol version 3 (L2TPv3) feature defines the L2TP protocol for tunneling Layer 2 traffic over an IP core network using Layer 2 VPNs. Benefits of this feature include:

- Simplifies deployment of VPNs
- Does not require Multiprotocol Label Switching (MPLS)
- Supports Layer 2 tunneling over IP for any traffic
- Supports data encapsulation directly over IP (IP protocol number 115), not using User Datagram Protocol (UDP)
- Supports point-to-point sessions, not point-to-multipoint or multipoint-to-point sessions
- Supports sessions between the same Layer 2 protocols, for example Frame Relay to Frame Relay or ATM to ATM

**Note**

---

This feature is supported on Cisco XR 12000 Series Routers.

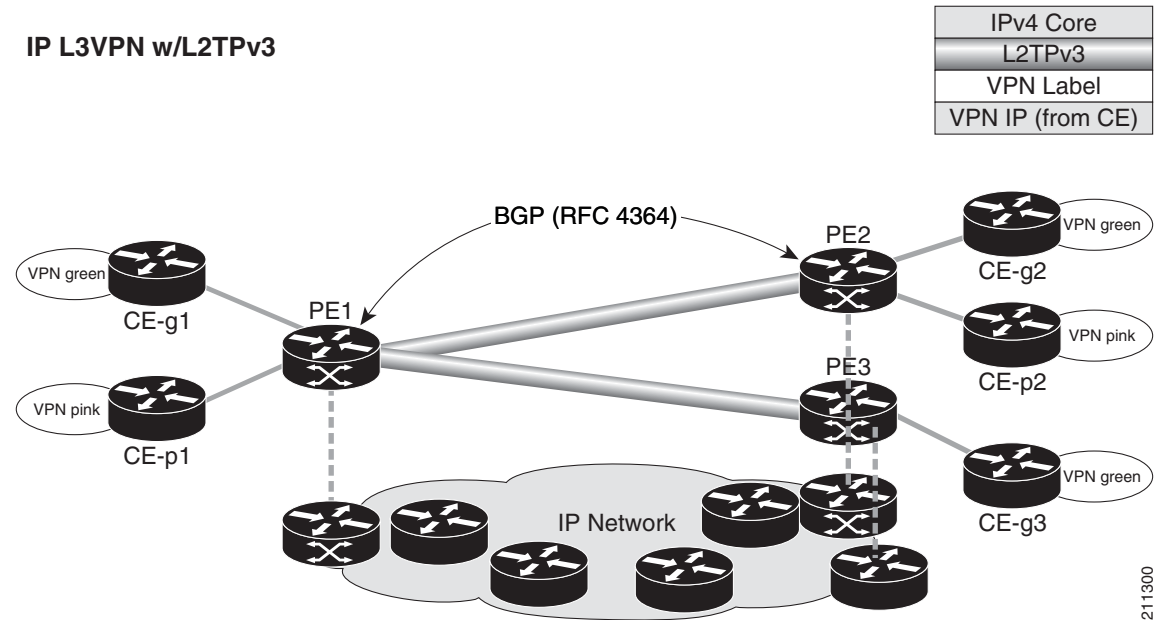
---

When an RFC 4364-based IP VPN service is deployed (see RFC 4364), VPN traffic is typically transported across the core network between service provider edge (PE) routers using MPLS label switched paths (LSPs). Native IP L3VPNs eliminate the need for MPLS between the participating core routers by relying on scalable tunnel encapsulation over IP. These tunnels can be used instead of, or with, MPLS to transport VPN traffic between participating edge routers.

A native IP L3VPN allows service providers to use an IP backbone to provide VPN services. BGP is used to distribute VPN routing information across the provider backbone.

[Figure 9](#) shows edge routers participating in switching IPv4 and IPv6 traffic over a tunnel using IP as the transport.

Figure 9 IP L3VPN with L2TPv3



211300

## BGP Multicast VPN

The BGP Multicast VPN feature introduces the IPv4 multicast distribution tree (MDT) subaddress family identifier (SAFI) in Border Gateway Protocol (BGP).

Multicast VPN (MVPN) extends the VPN architecture to provide multicast services over a shared service provider backbone using native multicast technology. This is achieved using virtual connections between provider edge (PE) routers in each VPN and using native multicast forwarding inside the provider network. An MDT may span across multiple customer sites and the provider network, allowing traffic to flow freely from one source to multiple receivers.

MVPN is supported on VPN networks based on MPLS and on networks based on IP Layer 2 Tunnel Protocol version 3 (L2TPv3).

PE routers are the only routers that must be MVPN-aware and that must be able to signal to remote PEs information regarding the MVPN. Therefore, all PE routers must have a BGP relationship with each other—either directly or using a route reflector (RR).

Generally the source address of the default MDT is the same address used to source the internal BGP (IBGP) sessions with the remote PE routers that belong to the same VPN and multicast VPN routing and forwarding (MVRF) instance. When Protocol Independent Multicast–Source Specific Multicast (PIM–SSM) is used for transport inside the provider core, it is through the BGP relationship that the PEs indicate that they are MVPN-capable and provide for source discovery. This capability is indicated using the updated BGP message.

When a PE receives a BGP update, which includes the rendezvous point (RP) and the group information, it joins the root of that tree, thereby joining the MDT.

Figure 10 shows Multiprotocol IBGP updates for MVPN. On PE1, PE2 is configured as its IBGP peer. This BGP peer configuration within a VRF triggers the MP-IBGP updates that send PE1 local VPN routes to PE2. BGP process on PE2 receives the VPN updates and installs VPN routes in the Routing Information Base (RIB) VRF table. When PIM looks up a VRF source or rendezvous point address that is reachable through the provider core, it receives an MP-IBGP route from the RIB.

When an MVPN-specific default MDT group is configured on PE1, PIM creates a virtual MDT tunnel interface with the tunnel source address the same as the BGP local peering address. This MDT interface is used by PIM to send VPN packets to the provider network and to receive VPN packets from the provider network. PIM also exchanges control messages over this MDT interface.

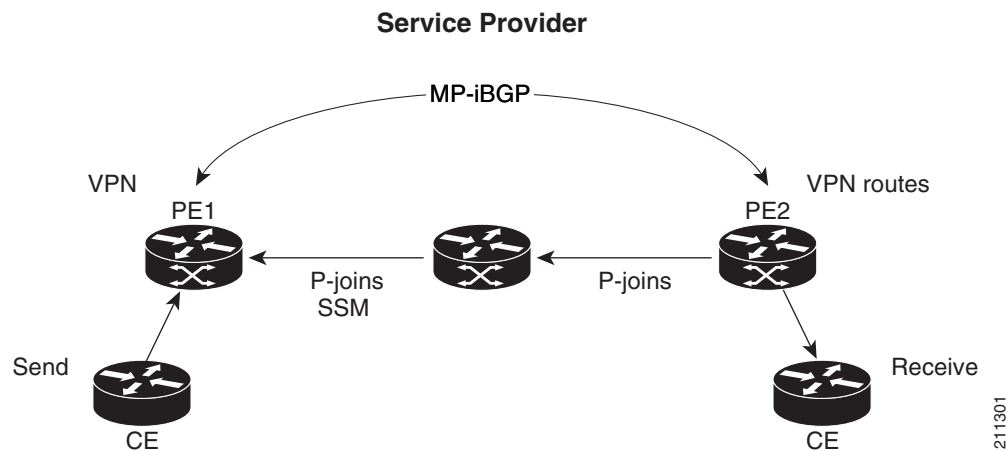
Each time a default MDT group is configured for a specific VRF, BGP builds an MDT SAFI update, with network layer reachability information (NLRI) containing the local PE BGP peering address and the newly configured MDT group address (The NLRI format is 8-byte-RD:IPv4-address followed by the MDT group address). This update is sent to all the BGP peers including PE2. The BGP process on PE2 receives this MDT update and notifies PIM. If the group is a PIM-SSM group, PIM on PE2 begins sending SSM joins to the BGP peering address on PE1 to establish an SSM tree in the core. This SSM tree is used to carry PIM control traffic and multicast data traffic in the corresponding VRF.

In summary, PIM requires the following from BGP:

- A new BGP MDT SAFI, which carries the VRF RD and BGP local peering address and default MDT group in its NLRI.
  - A notification mechanism from BGP to PIM about the availability of the MDT SAFI update.
  - A notification mechanism from PIM to BGP about the default MDT group address and source address.

See Internet Engineering Task Force (IETF) draft-nalawade-idr-mdt-safi-03 for detailed information on MDT SAFI.

**Figure 10** Multiprotocol IBGP Updates for MVPN



## How to Implement BGP on Cisco IOS XR Software

This section contains instructions for the following tasks:

- [Enabling BGP Routing, page RC-44](#) (required)
- [Configuring a Routing Domain Confederation for BGP, page RC-47](#) (optional)
- [Resetting an eBGP Session Immediately Upon Link Failure, page RC-49](#) (optional)
- [Logging Neighbor Changes, page RC-49](#) (optional)
- [Adjusting BGP Timers, page RC-50](#) (optional)
- [Changing the BGP Default Local Preference Value, page RC-51](#) (optional)



- [Configuring the MED Metric for BGP, page RC-53](#) (optional)
- [Configuring BGP Weights, page RC-54](#) (optional)
- [Tuning the BGP Best-Path Calculation, page RC-56](#) (optional)
- [Indicating BGP Back-door Routes, page RC-58](#) (optional)
- [Configuring Aggregate Addresses, page RC-60](#) (optional)
- [Redistributing iBGP Routes into IGP, page RC-61](#) (optional)
- [Redistributing Prefixes into Multiprotocol BGP, page RC-63](#) (optional)
- [Configuring BGP Route Dampening, page RC-66](#) (optional)
- [Applying Policy When Updating the Routing Table, page RC-71](#) (optional)
- [Setting BGP Administrative Distance, page RC-73](#) (optional)
- [Configuring a BGP Neighbor Group and Neighbors, page RC-74](#) (required)
- [Configuring a Route Reflector for BGP, page RC-77](#) (optional)
- [Configuring BGP Route Filtering by Route Policy, page RC-79](#) (optional)
- [Configuring BGP Next-hop Trigger Delay, page RC-81](#) (optional)
- [Disabling Next-hop Processing on BGP Updates, page RC-83](#) (optional)
- [Configuring BGP Community and Extended-Community Advertisements, page RC-84](#) (optional)
- [Configuring the BGP Cost Community, page RC-86](#) (optional)
- [Configuring Software to Store Updates from a Neighbor, page RC-91](#) (optional)
- [Configuring Distributed BGP, page RC-93](#) (optional)
- [Configuring a VPN Routing and Forwarding Instance in BGP, page RC-96](#) (optional)
- [Configuring Keychains for BGP, page RC-114](#) (optional)
- [Configuring an MDT Address Family Session in BGP, page RC-115](#) (optional)
- [Disabling a BGP Neighbor, page RC-118](#) (optional)
- [Resetting Neighbors Using BGP Inbound Soft Reset, page RC-120](#) (optional)
- [Resetting Neighbors Using BGP Outbound Soft Reset, page RC-120](#) (optional)
- [Resetting Neighbors Using BGP Hard Reset, page RC-121](#) (optional)
- [Clearing Caches, Tables, and Databases, page RC-122](#) (optional)
- [Displaying System and Network Statistics, page RC-123](#) (optional)
- [Displaying BGP Process Information, page RC-125](#) (optional)
- [Monitoring BGP Update Groups, page RC-126](#) (optional)

## Enabling BGP Routing

Perform this task to enable BGP routing and establish a BGP routing process. Configuring BGP neighbors is included as part of enabling BGP routing.



### Note

At least one neighbor and at least one address family must be configured to enable BGP routing. At least one neighbor with both a remote AS and an address family must be configured globally using the **address family** and **remote as** commands.

## Prerequisites

BGP must be able to obtain a router identifier (for example, a configured loopback address). At least, one address family must be configured in the BGP router configuration and the same address family must also be configured under the neighbor.

## Restrictions

If the neighbor is configured as an external BGP (eBGP) peer, you must configure an inbound and outbound route policy on the neighbor using the **route-policy** command.

## SUMMARY STEPS

1. **configure**
2. **route-policy** *route-policy-name*
3. **end-policy**
4. **end**  
or  
**commit**
5. **configure**
6. **router bgp** *as-number*
7. **bgp router-id** *ip-address*
8. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast** }
9. **exit**
10. **neighbor** *ip-address*
11. **remote-as** *as-number*
12. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast** }
13. **route-policy** *route-policy-name* { **in** | **out** }
14. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mod.
Step 2	<p><b>route-policy</b> <i>route-policy-name</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy drop-as-1234 RP/0/RP0/CPU0:router(config-rpl)# if as-path passes-through '1234' then RP/0/RP0/CPU0:router(config-rpl)# apply check-communities RP/0/RP0/CPU0:router(config-rpl)# else RP/0/RP0/CPU0:router(config-rpl)# pass RP/0/RP0/CPU0:router(config-rpl)# endif</p>	(Optional) Defines a route policy and enters route policy configuration mode.
Step 3	<p><b>end-policy</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end-policy</p>	(Optional) Ends the definition of a route policy and exits route policy configuration mode.
Step 4	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# end OR RP/0/RP0/CPU0:router(config)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

	Command or Action	Purpose
Step 5	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 6	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 7	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp router-id 192.168.70.24	Configures the local router with a specified router ID.
Step 8	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>vpnv4 unicast</i>   <i>vpnv6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters global address family configuration mode for the specified address family.
Step 9	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 10	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 11	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
Step 12	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 labeled-unicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>ipv6 labeled-unicast</i>   <i>vpnv4 unicast</i>   <i>vpnv6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 13	<pre>route-policy route-policy-name {in   out}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)#  route-policy drop-as-1234 in </p>	(Optional) Applies the specified policy to inbound IPv4 unicast routes.
Step 14	<pre>end</pre> <p>or</p> <pre>commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end</p> <p>or</p> <pre>RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?</pre> <pre>[cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a Routing Domain Confederation for BGP

Perform this task to configure the routing domain confederation for BGP. This includes specifying a confederation identifier and autonomous systems that belong to the confederation.

Configuring a routing domain confederation reduces the internal BGP (iBGP) mesh by dividing an autonomous system into multiple autonomous systems and grouping them into a single confederation. Each autonomous system is fully meshed within itself and has a few connections to another autonomous system in the same confederation. The confederation maintains the next hop and local preference information, and that allows you to retain a single Interior Gateway Protocol (IGP) for all autonomous systems. To the outside world, the confederation looks like a single autonomous system.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp confederation identifier** *as-number*
4. **bgp confederation peers** *as-number*
5. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>bgp confederation identifier</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp confederation identifier 5	Specifies a BGP confederation identifier.

	Command or Action	Purpose
Step 4	<p><b>bgp confederation peers</b> <i>as-number</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1091  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1092  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1093  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1094  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1095  RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1096</p>	<p>Specifies that the BGP autonomous systems belong to a specified BGP confederation identifier.</p>
Step 5	<p><b>end</b>  OR  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp)# end  OR  RP/0/RP0/CPU0:router(config-bgp)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)?  [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Resetting an eBGP Session Immediately Upon Link Failure

Immediately resetting BGP sessions of any directly adjacent external peers if the link used to reach them goes down is enabled by default. Use the **bgp fast-external-falover disable** command to disable automatic resetting. The **no bgp fast-external-falover disable** command can also be used to turn the automatic reset back on.

## Logging Neighbor Changes

Logging neighbor changes is enabled by default. Use the **log neighbor changes disable** command to turn off logging. The **no log neighbor changes disable** command can also be used to turn logging back on if it has been disabled.

## Adjusting BGP Timers

Perform this task to set the timers for BGP neighbors.

BGP uses certain timers to control periodic activities, such as the sending of keepalive messages and the interval after which a neighbor is assumed to be down if no messages are received from the neighbor during the interval. The values set using the **timers bgp** command in router configuration mode can be overridden on particular neighbors using the **timers** command in the neighbor configuration mode.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **timers bgp** *keepalive hold-time*
4. **neighbor** *ip-address*
5. **timers** *keepalive hold-time*
6. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 123	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>timers bgp</b> <i>keepalive hold-time</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# timers bgp 30 90	Sets a default keepalive time and a default hold time for all neighbors.
Step 4	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.



	Command or Action	Purpose
Step 5	<p><b>timers</b> <i>keepalive hold-time</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 60 220</p>	(Optional) Sets the keepalive timer and the hold-time timer for the BGP neighbor.
Step 6	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# end or RP/0/RP0/CPU0:router(config-bgp-nbr)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Changing the BGP Default Local Preference Value

Perform this task to set the default local preference value for BGP paths.

### SUMMARY STEPS

- configure**
- router** **bgp** *as-number*
- bgp** **default** **local-preference** *value*
- end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120</p>	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<p><b>bgp default local-preference</b> <i>value</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp default local-preference 200</p>	Sets the default local preference value from the default of 100, making it either a more preferable path (over 100) or less preferable path (under 100).
Step 4	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# end OR RP/0/RP0/CPU0:router(config-bgp)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring the MED Metric for BGP

Perform this task to set the multi exit discriminator (MED) to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **default-metric** *value*
4. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.

Command or Action	Purpose
<p><b>Step 3</b></p> <p><b>default-metric</b> <i>value</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp)# default  metric 10</p>	<p>Sets the default metric, which is used to set the MED to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).</p>
<p><b>Step 4</b></p> <p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp)# end  OR  RP/0/RP0/CPU0:router(config-bgp)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Weights

Perform this task to assign a weight to routes received from a neighbor. A weight is a number that you can assign to a path so that you can control the best-path selection process. If you have particular neighbors that you want to prefer for most of your traffic, you can use the **weight** command to assign a higher weight to all routes learned from that neighbor.

### Restrictions

The **clear bgp** command must be used for the newly configured weight to take effect.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpnv4 unicast** | **vpn6 unicast** }
6. **weight** *weight-value*

```

7. end
   or
   commit
    
```

**DETAILED STEPS**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
<b>Step 2</b>	<p><b>router bgp as-number</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120</p>	Enters BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>neighbor ip-address</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24</p>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<p><b>remote-as as-number</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002</p>	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 5</b>	<p><b>address-family {ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpnv4 unicast   vpnv6 unicast}</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast</p>	Enters neighbor address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 6	<p><b>weight</b> <i>weight-value</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# weight 41150</p>	Assigns a weight to all routes learned through the neighbor.
Step 7	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end or RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</li> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Tuning the BGP Best-Path Calculation

Perform this task to change the default BGP best-path calculation behavior.

### SUMMARY STEPS

- configure**
- router** **bgp** *as-number*
- bgp bestpath med missing-as-worst**
- bgp bestpath med always**
- bgp bestpath med confed**
- bgp bestpath as-path ignore**
- bgp bestpath compare-routerid**
- end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp as-number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 126	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>bgp bestpath med missing-as-worst</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst	Directs the BGP software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path.
Step 4	<b>bgp bestpath med always</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med always	Configures the BGP speaker in the specified autonomous system to compare MEDs among all the paths for the prefix, regardless of the autonomous system from which the paths are received.
Step 5	<b>bgp bestpath med confed</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med confed	Enables BGP software to compare MED values for paths learned from confederation peers.
Step 6	<b>bgp bestpath as-path ignore</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath as-path ignore	Configures the BGP software to ignore the autonomous system length when performing best-path selection.

	Command or Action	Purpose
Step 7	<pre> <b>bgp bestpath compare-routerid</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath compare-routerid </pre>	Configure the BGP speaker in the autonomous system to compare the router IDs of similar paths.
Step 8	<pre> <b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# end or RP/0/RP0/CPU0:router(config-bgp)# commit </pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: </pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Indicating BGP Back-door Routes

Perform this task to set the administrative distance on an external Border Gateway Protocol (eBGP) route to that of a locally sourced BGP route, causing it to be less preferred than an Interior Gateway Protocol (IGP) route.

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv6 unicast* | *ipv6 multicast* }
- network** { *ip-address /prefix-length* | *ip-address mask* } **backdoor**
- end**  
or  
**commit**



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120</p>	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<p><b>address-family</b> {<i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast</p>	Enters address family configuration mode for the specified address family.
Step 4	<p><b>network</b> {<i>ip-address /prefix-length</i>   <i>ip-address mask</i>} <b>backdoor</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# network 172.20.0.0/16</p>	Configures the local router to originate and advertise the specified network.
Step 5	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# end OR RP/0/RP0/CPU0:router(config-bgp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Aggregate Addresses

Perform this task to create aggregate entries in a BGP routing table.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv6 unicast** | **ipv6 multicast** }
4. **aggregate-address** *address/mask-length* [**as-set**] [**as-confed-set**] [**summary-only**] [**route-policy** *route-policy-name*]
5. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <b>ipv4 unicast</b>   <b>ipv4 multicast</b>     <b>ipv6 unicast</b>   <b>ipv6 multicast</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.

Command or Action	Purpose
<p><b>Step 4</b></p> <pre>aggregate-address address/mask-length [as-set] [as-confed-set] [summary-only] [route-policy route-policy-name]</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set</pre>	<p>Creates an aggregate address. The path advertised for this route is an autonomous system set consisting of all elements contained in all paths that are being summarized.</p> <ul style="list-style-type: none"> <li>• The <b>as-set</b> keyword generates autonomous system set path information and community information from contributing paths.</li> <li>• The <b>as-confed-set</b> keyword generates autonomous system confederation set path information from contributing paths.</li> <li>• The <b>summary-only</b> keyword filters all more specific routes from updates.</li> <li>• The <b>route-policy route-policy-name</b> keyword and argument specify the route policy used to set the attributes of the aggregate route.</li> </ul>
<p><b>Step 5</b></p> <pre>end OR commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-bgp-af)# end OR RP/0/RP0/CPU0:router(config-bgp-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing iBGP Routes into IGP

Perform this task to redistribute iBGP routes into an Interior Gateway Protocol (IGP), such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF).



### Note

Use of the **bgp redistribute-internal** command requires the **clear route \*** command to be issued to reinstall all BGP routes into the IP routing table.



### Caution

Redistributing iBGP routes into IGPs may cause routing loops to form within an autonomous system. Use this command with caution.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp redistribute-internal**
4. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>bgp redistribute-internal</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp redistribute-internal	Allows the redistribution of iBGP routes into an IGP, such as IS-IS or OSPF.
Step 4	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# end or RP/0/RP0/CPU0:router(config-bgp)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing Prefixes into Multiprotocol BGP

Perform this task to redistribute prefixes from another protocol into multiprotocol BGP.

Redistribution is the process of injecting prefixes from one routing protocol into another routing protocol. This task shows how to inject prefixes from another routing protocol into multiprotocol BGP. Specifically, prefixes that are redistributed into multiprotocol BGP using the **redistribute** command are injected into the unicast database, the multicast database, or both.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv6 unicast* | *ipv6 multicast* }
4. **redistribute connected** [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute eigrp** *process-id* [**match** { *external* | *internal* }] [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute isis** *process-id* [**level** { *1* | *1-inter-area* | *2* }] [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute ospf** *process-id* [**match** { *external* [*1* | *2*] | *internal* | *nssa-external* [*1* | *2*]}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute ospfv3** *process-id* [**match** { *external* [*1* | *2*] | *internal* | *nssa-external* [*1* | *2*]}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute rip** [**metric** *metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute static** [**metric** *metric-value*] [**route-policy** *route-policy-name*]
5. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.

Command or Action	Purpose
<p><b>Step 3</b></p> <pre>address-family {ipv4 unicast   ipv4 multicast     ipv6 unicast   ipv6 multicast}</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast</pre>	Enters address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 4	<pre> <b>redistribute connected</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute eigrp</b> <i>process-id</i> [<b>match</b> {<b>external</b>   <b>internal</b>}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute isis</b> <i>process-id</i> [<b>level</b> {<b>1</b>   <b>1-inter-area</b>   <b>2</b>}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute ospf</b> <i>process-id</i> [<b>match</b> {<b>external</b> [<b>1</b>   <b>2</b>]   <b>internal</b>   <b>nssa-external</b> [<b>1</b>   <b>2</b>]}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute ospfv3</b> <i>process-id</i> [<b>match</b> {<b>external</b> [<b>1</b>   <b>2</b>]   <b>internal</b>   <b>nssa-external</b> [<b>1</b>   <b>2</b>]}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute rip</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute static</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>]  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# redistribute ospf 110 </pre>	<p>Causes routes from the specified instance to be redistributed into BGP.</p>
Step 5	<pre> <b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# <b>end</b> OR RP/0/RP0/CPU0:router(config-bgp-af)# <b>commit</b> </pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: </pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Route Dampening

Perform this task to configure and monitor BGP route dampening.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv6 unicast** | **ipv6 multicast**}
4. **bgp dampening** [*half-life* [*reuse suppress max-suppress-time*] | **route-policy** *route-policy-name*]
5. **end**  
or  
**commit**
6. **show bgp** [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpn6 unicast** [**rd** *rd-address*]] **flap-statistics**
7. **show bgp** [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpn6 unicast** [**rd** *rd-address*]] **flap-statistics** **regex** *regular-expression*
8. **show bgp** [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpn6 unicast** [**rd** *rd-address*]] **flap-statistics** **route-policy** *route-policy-name*
9. **show bgp** [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpn6 unicast** [**rd** *rd-address*]] **flap-statistics** {*ip-address* {*mask* | */prefix-length*}}
10. **show bgp** [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpn6 unicast** [**rd** *rd-address*]] **flap-statistics** {*ip-address* [{*mask* | */prefix-length*] [**longer-prefixes**]}
11. **clear bgp** {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** | **vrf** {*vrf-name* | **all**} | **ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**} | **vpn6 unicast**} **flap-statistics**
12. **clear bgp** {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** | **vrf** {*vrf-name* | **all**} | **ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**} | **vpn6 unicast**} **flap-statistics** **regex** *regular-expression*
13. **clear bgp** {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **vpn4 unicast** | **vrf** {*vrf-name* | **all**} | **ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**} | **vpn6 unicast**} **flap-statistics** **route-policy** *route-policy-name*



14. **clear bgp** {ipv4 {unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} flap-statistics network/mask-length
15. **clear bgp** {ipv4 {unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} flap-statistics ip-address/mask-length
16. **show bgp** [ipv4 {unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast} | vpnv4 unicast [rd rd-address] | vrf {vrf-name | all} [ipv4 {unicast | labeled-unicast} | ipv6 unicast] | vpnv6 unicast [rd rd-address]] dampened-paths
17. **clear bgp** {ipv4 {unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} dampening [ip-address/mask-length]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> as-number  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> {ipv4 unicast   ipv4 multicast   ipv6 unicast   ipv6 multicast}  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.
Step 4	<b>bgp dampening</b> [half-life [reuse suppress max-suppress-time]   route-policy route-policy-name]  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# bgp dampening 30 1500 10000 120	Configures BGP dampening for the specified address family.

	Command or Action	Purpose
Step 5	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)# end  OR  RP/0/RP0/CPU0:router(config-bgp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 6	<pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   tunnel}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast [rd rd-address]] flap-statistics</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show bgp flap statistics</p>	<p>Displays BGP flap statistics.</p>
Step 7	<pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast [rd rd-address]] flap-statistics regexp regular-expression</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show bgp flap-statistics regexp _1\$</p>	<p>Displays BGP flap statistics for all paths that match the regular expression.</p>

	Command or Action	Purpose
<p><b>Step 8</b></p>	<pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast [rd rd-address]] flap-statistics route-policy route-policy-name</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config)# show bgp flap-statistics route-policy policy_A</p>	<p>Displays BGP flap statistics for the specified route policy.</p>
<p><b>Step 9</b></p>	<pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast [rd rd-address]] flap-statistics {ip-address {mask   /prefix-length}}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show bgp flap-statistics 172.20.1.1</p>	<p>Displays BGP flap for the specified prefix.</p>
<p><b>Step 10</b></p>	<pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast [rd rd-address]] flap-statistics {ip-address [{mask   /prefix-length} [longer-prefixes]]}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show bgp flap-statistics 172.20.1.1 longer-prefixes</p>	<p>Displays BGP flap statistics for more specific entries for the specified IP address.</p>
<p><b>Step 11</b></p>	<pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpngv6 unicast} flap-statistics</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# clear bgp all all flap-statistics</p>	<p>Clears BGP flap statistics for all routes.</p>

Command or Action	Purpose
<p><b>Step 12</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} flap-statistics regexp regular-expression</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics regexp _1\$</p>	<p>Clears BGP flap statistics for all paths that match the specified regular expression.</p>
<p><b>Step 13</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} flap-statistics route-policy route-policy-name</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics route-policy policy_A</p>	<p>Clears BGP flap statistics for the specified route policy.</p>
<p><b>Step 14</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} flap-statistics network/mask-length</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics 192.168.40.0/24</p>	<p>Clears BGP flap statistics for the specified network.</p>
<p><b>Step 15</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} flap-statistics ip-address/mask-length</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics 172.20.1.1</p>	<p>Clears BGP flap statistics for routes received from the specified neighbor.</p>

Command or Action	Purpose
<p><b>Step 16</b></p> <pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpn4 unicast [rd rd-address]   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpn6 unicast [rd rd-address]] dampened-paths</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp dampened paths</p>	<p>Displays the dampened routes, including the time remaining before they are unsuppressed.</p>
<p><b>Step 17</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast}   vpn4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpn6 unicast} dampening [ip-address/mask-length]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp dampening</p>	<p>Clears route dampening information and unsuppresses the suppressed routes.</p>

## Applying Policy When Updating the Routing Table

Perform this task to apply a routing policy to routes being installed into the routing table.

### Prerequisites

See the *Implementing Routing Policy on Cisco IOS XR Software* module of *Cisco IOS XR Routing Configuration Guide* for a list of the supported attributes and operations that are valid for table policy filtering.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** {*ipv4 unicast* | *ipv4 multicast* | *ipv6 unicast* | *ipv6 multicast*}
4. **table-policy** *policy-name*
5. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120.6	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>     <i>ipv6 unicast</i>   <i>ipv6 multicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.
Step 4	<b>table-policy</b> <i>policy-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# table-policy tbl-plcy-A	Applies the specified policy to routes being installed into the routing table.
Step 5	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# end or RP/0/RP0/CPU0:router(config-bgp-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Setting BGP Administrative Distance

Perform this task to specify the use of administrative distances that can be used to prefer one class of route over another.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** {*ipv4 unicast* | *ipv4 multicast* | *ipv6 unicast* | *ipv6 multicast*}
4. **distance bgp** *external-distance internal-distance local-distance*
5. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>     <i>ipv6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 4	<pre>distance bgp external-distance internal-distance local-distance</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)# distance  bgp 20 20 200 </p>	Sets the external, internal, and local administrative distances to prefer one class of routes over another. The higher the value, the lower the trust rating.
Step 5	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)# end  or  RP/0/RP0/CPU0:router(config-bgp-af)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a BGP Neighbor Group and Neighbors

Perform this task to configure BGP neighbor groups and apply the neighbor group configuration to a neighbor. A neighbor group is a template that holds address family-independent and address family-dependent configurations associated with the neighbor.

After a neighbor group is configured, each neighbor can inherit the configuration through the **use** command. If a neighbor is configured to use a neighbor group, the neighbor (by default) inherits the entire configuration of the neighbor group, which includes the address family-independent and address family-dependent configurations. The inherited configuration can be overridden if you directly configure commands for the neighbor or configure session groups or address family groups through the **use** command.

You can configure an address family-independent configuration under the neighbor group. An address family-dependent configuration requires you to configure the address family under the neighbor group to enter address family submode.

From neighbor group configuration mode, you can configure address family-independent parameters for the neighbor group. Use the **address-family** command when in the neighbor group configuration mode.

After specifying the neighbor group name using the **neighbor group** command, you can assign options to the neighbor group.





**Note**

All commands that can be configured under a specified neighbor group can be configured under a neighbor.

**SUMMARY STEPS**

1. **configure**
2. **router bgp** *as-number*
3. **address-family** {*ipv4 unicast | ipv4 multicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6 multicast | vpnv4 unicast | vpnv6 unicast*}
4. **exit**
5. **neighbor-group** *name*
6. **remote-as** *as-number*
7. **address-family** {*ipv4 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6 multicast | ipv6 labeled-unicast | vpnv4 unicast | vpnv6 unicast*}
8. **route-policy** *route-policy-name* {*in | out*}
9. **exit**
10. **exit**
11. **neighbor** *ip-address*
12. **use neighbor-group** *group-name*
13. **remote-as** *as-number*
14. **end**  
or  
**commit**

**DETAILED STEPS**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>aa-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <i>ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast</i> }	Enters neighbor group address family configuration mode for the specified address family.
	<b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	

	Command or Action	Purpose
Step 4	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 5	<b>neighbor-group</b> <i>name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor-group nbr-grp-A	Places the router in neighbor group configuration mode.
Step 6	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
Step 7	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 labeled-unicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>ipv6 labeled-unicast</i>   <i>vpn4 unicast</i>   <i>vpn6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast	Enters neighbor group address family configuration mode for the specified address family.
Step 8	<b>route-policy</b> <i>route-policy-name</i> { <i>in</i>   <i>out</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy drop-as-1234 in	(Optional) Applies the specified policy to inbound IPv4 unicast routes.
Step 9	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit	Exits the current configuration mode.
Step 10	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit	Exits the current configuration mode.
Step 11	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 12	<b>use neighbor-group</b> <i>group-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group nbr-grp-A	(Optional) Specifies that the BGP neighbor inherit configuration from the specified neighbor group.

	Command or Action	Purpose
Step 13	<pre>remote-as as-number</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002</p>	Creates a neighbor and assigns a remote autonomous system number to it.
Step 14	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a Route Reflector for BGP

Perform this task to configure a route reflector for BGP.

All the neighbors configured with the **route-reflector-client** command are members of the client group, and the remaining iBGP peers are members of the nonclient group for the local route reflector.

Together, a route reflector and its clients form a *cluster*. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the software as the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, a cluster can have more than one route reflector. If it does, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. The **bgp cluster-id** command is used to configure the cluster ID when the cluster has more than one route reflector.

### SUMMARY STEPS

1. **configure**
2. **router bgp as-number**
3. **bgp cluster-id cluster-id**
4. **neighbor ip-address**
5. **remote-as as-number**

6. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast** }
7. **route-reflector-client**
8. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>bgp cluster-id</b> <i>cluster-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 192.168.70.1	Configures the local router as one of the route reflectors serving the cluster. It is configured with a specified cluster ID to identify the cluster.
Step 4	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 5	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2003	Creates a neighbor and assigns a remote autonomous system number to it.
Step 6	<b>address-family</b> { <b>ipv4 unicast</b>   <b>ipv4 multicast</b>   <b>ipv4 labeled-unicast</b>   <b>ipv4 tunnel</b>   <b>ipv4 mdt</b>   <b>ipv6 unicast</b>   <b>ipv6 multicast</b>   <b>ipv6 labeled-unicast</b>   <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast	Enters neighbor address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 7	<pre>route-reflector-client</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client </p>	Configures the router as a BGP route reflector and configures the neighbor as its client.
Step 8	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Route Filtering by Route Policy

Perform this task to configure BGP routing filtering by route policy.

### Prerequisites

See the *Implementing Routing Policy on Cisco IOS XR Software* module of *Cisco IOS XR Routing Configuration Guide* for a list of the supported attributes and operations that are valid for inbound and outbound neighbor policy filtering.

### SUMMARY STEPS

- configure**
- route-policy** *name*
- end-policy**
- router bgp** *as-number*
- neighbor** *ip-address*
- address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast** }
- route-policy** *route-policy-name* { **in** | **out** }

```

8. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy drop-as-1234 RP/0/RP0/CPU0:router(config-rpl)# if as-path passes-through '1234' then RP/0/RP0/CPU0:router(config-rpl)# apply check-communities RP/0/RP0/CPU0:router(config-rpl)# else RP/0/RP0/CPU0:router(config-rpl)# pass RP/0/RP0/CPU0:router(config-rpl)# endif	(Optional) Defines a route policy and enters route policy configuration mode.
Step 3	<b>end-policy</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end-policy	(Optional) Ends the definition of a route policy and exits route policy configuration mode.
Step 4	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 5	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 6	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 labeled-unicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>ipv6 labeled-unicast</i>   <i>vpn4 unicast</i>   <i>vpn6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast	Enters neighbor address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 7	<pre>route-policy route-policy-name {in   out}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy drop-as-1234 in </p>	Applies the specified policy to inbound routes.
Step 8	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Next-hop Trigger Delay

Perform this task to configure BGP next-hop trigger delay. The Routing Information Base (RIB) classifies the dampening notifications based on the severity of the changes. Event notifications are classified as critical and noncritical. This task allows you to specify the minimum batching interval for the critical and noncritical events.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** {*ipv4 unicast | ipv4 multicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6 multicast | vpnv4 unicast | vpnv6 unicast*}
4. **nexthop trigger-delay** {*critical delay | non-critical delay*}
5. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>vpnv4 unicast</i>   <i>vpnv6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family <i>ipv4 unicast</i>	Enters address family configuration mode for the specified address family.
Step 4	<b>nexthop trigger-delay</b> { <i>critical delay</i> / <i>non-critical delay</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# nexthop trigger-delay <i>critical 15000</i>	Sets the critical next-hop trigger delay.
Step 5	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# end OR RP/0/RP0/CPU0:router(config-bgp-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>



## Disabling Next-hop Processing on BGP Updates

Perform this task to disable next-hop calculation for a neighbor and insert your own address in the next-hop field of BGP updates. Disabling the calculation of the best next hop to use when advertising a route causes all routes to be advertised with the network device as the next hop.



### Note

Next-hop processing can be disabled for address family group, neighbor group, or neighbor address family.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family**{ **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast**}
6. **next-hop-self**
7. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 4	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 206	Creates a neighbor and assigns a remote autonomous system number to it.

	Command or Action	Purpose
Step 5	<pre>address-family {ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpnv4 unicast   vpnv6 unicast}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast</p>	Enters neighbor address family configuration mode for the specified address family.
Step 6	<pre>next-hop-self</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self</p>	Sets the next-hop attribute for all routes advertised to the specified neighbor to the address of the local router. Disabling the calculation of the best next hop to use when advertising a route causes all routes to be advertised with the local network device as the next hop.
Step 7	<pre>end OR commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Community and Extended-Community Advertisements

Perform this task to specify that community attributes should be sent to an eBGP neighbor.

Perform this task to specify that community/extended-community attributes should be sent to an eBGP neighbor. These attributes are not sent to an eBGP neighbor by default. By contrast, they are always sent to iBGP neighbors. This section provides examples on how to enable sending community attributes. The **send-community-ebgp** keyword can be replaced by the **send-extended-community-ebgp** keyword to enable sending extended-communities.



### Note

If the **send-community-ebgp** command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.

**Note**

BGP community and extended-community filtering cannot be configured for iBGP neighbors. Communities and extended-communities are always sent to iBGP neighbors

**SUMMARY STEPS**

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast** }
6. **send-community-ebgp**
7. **end**  
or  
**commit**

**DETAILED STEPS**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.

	Command or Action	Purpose
Step 5	<pre>address-family {ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpnv4 unicast   vpnv6 unicast}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)#  address-family ipv4 unicast</p>	Enters neighbor address family configuration mode for the specified address family.
Step 6	<pre>send-community-ebgp</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)#  send-community-ebgp</p>	Specifies that the router send community attributes (which are disabled by default for eBGP neighbors) to a specified eBGP neighbor.
Step 7	<pre>end OR commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end  OR  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring the BGP Cost Community

Perform this task to configure the BGP cost community.

BGP receives multiple paths to the same destination and it uses the best-path algorithm to decide which is the best path to install in RIB. To enable users to determine an exit point after partial comparison, the cost community is defined to tie-break equal paths during the best-path selection process.

### SUMMARY STEPS

1. **configure**
2. **route-policy** *name*
3. **set extcommunity cost** { *cost-extcommunity-set-name* | *cost-inline-extcommunity-set* } [**additive**]
4. **end-policy**

**5. router bgp** *as-number***6. default-information originate**

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**aggregate-address** *address/mask-length* [**as-set**] [**as-confed-set**] [**summary-only**] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute connected** [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute eigrp** *process-id* [**match** {**external** | **internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute isis** *process-id* [**level** {**1** | **1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute ospf** *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute ospfv3** *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

**address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast**}**redistribute rip** [**metric** *metric-value*] [**route-policy** *route-policy-name*]

or

```
address-family {ipv4 unicast | ipv4 multicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6
multicast | vpnv4 unicast | vpnv6 unicast}
```

```
redistribute static [metric metric-value] [route-policy route-policy-name]
```

or

```
address-family {ipv4 unicast | ipv4 multicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6
multicast | vpnv4 unicast | vpnv6 unicast}
```

```
network {ip-address/prefix-length | ip-address mask} [route-policy route-policy-name]
```

or

```
neighbor ip-address
```

```
remote-as as-number
```

```
address-family {ipv4 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | ipv4 mdt |
ipv6 unicast | ipv6 multicast | ipv6 labeled-unicast | vpnv4 unicast | vpnv6 unicast}
```

```
route-policy route-policy-name {in | out}
```

7. **end**  
or  
**commit**
8. **show bgp** [**vrf** *vrf-name*] *ip-address*

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy costA	Enters route policy configuration mode and specifies the name of the route policy to be configured.
Step 3	<b>set extcommunity</b> <i>cost</i> { <i>cost-extcommunity-set-name</i>   <i>cost-inline-extcommunity-set</i> } [ <b>additive</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config)# set extcommunity cost cost_A	Specifies the BGP extended community attribute for cost.
Step 4	<b>end-policy</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# end-policy	Ends the definition of a route policy and exits route policy configuration mode.

	Command or Action	Purpose
Step 5	<pre>router bgp as-number</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120</p>	<p>Enters BGP configuration mode allowing you to configure the BGP routing process.</p>
Step 6	<pre>default-information originate or aggregate-address address/mask-length [as-set] [as-confed-set] [summary-only] [route-policy route-policy-name] or address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast}  redistribute connected [metric metric-value] [route-policy route-policy-name] or address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast}  redistribute eigrp process-id [match {external   internal}] [metric metric-value] [route-policy route-policy-name] or address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast}  redistribute isis process-id [level {1   1-inter-area   2}] [metric metric-value] [route-policy route-policy-name] or address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast}  redistribute ospf process-id [match {external [1   2]   internal   nssa-external [1   2]}] [metric metric-value] [route-policy route-policy-name] or</pre>	<p>Applies the cost community to the attach point (route policy).</p>

Command or Action	Purpose
<pre> address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast}  redistribute ospfv3 process-id [match {external [1   2]   internal   nssa-external [1   2]]} [metric metric-value] [route-policy route-policy-name] OR address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast}  redistribute rip [metric metric-value] [route-policy route-policy-name] OR address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast}  redistribute static [metric metric-value] [route-policy route-policy-name] OR address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast}  network {ip-address/prefix-length   ip-address mask} [route-policy route-policy-name] OR neighbor ip-address  remote-as as-number  address-family {ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpnv4 unicast   vpnv6 unicast}  route-policy route-policy-name {in   out} OR </pre>	
<p><b>Example:</b></p> <pre> RP/0/RP0/CPU0:router(config-bgp)# default-information originate </pre>	



	Command or Action	Purpose
Step 7	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)# end  OR  RP/0/RP0/CPU0:router(config-bgp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 8	<pre>show bgp [vrf vrf-name] ip-address</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show bgp 172.168.40.24</p>	<p>Displays the cost community in the following format:  Cost:&lt;POI&gt;:&lt;cost community ID&gt;:&lt;cost number&gt;</p>

## Configuring Software to Store Updates from a Neighbor

Perform this task to configure the software to store updates received from a neighbor.

The **soft-reconfiguration inbound** command causes a route refresh request to be sent to the neighbor if the neighbor is route refresh capable. If the neighbor is not route refresh capable, the neighbor must be reset to relearn received routes using the **clear bgp soft** command. See the “[Resetting Neighbors Using BGP Inbound Soft Reset](#)” section on page RC-120.



### Note

Storing updates from a neighbor works only if either the neighbor is route refresh capable or the **soft-reconfiguration inbound** command is configured. Even if the neighbor is route refresh capable and the **soft-reconfiguration inbound** command is configured, the original routes are not stored unless the **always** option is used with the command. The original routes can be easily retrieved with a route refresh request. Route refresh sends a request to the peer to resend its routing information. The **soft-reconfiguration inbound** command stores all paths received from the peer in an unmodified form and refers to these stored paths during the clear. Soft reconfiguration is memory intensive.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*

4. **address-family** {**ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast**}
5. **soft-reconfiguration inbound** [**always**]
6. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 4	<b>address-family</b> { <b>ipv4 unicast</b>   <b>ipv4 multicast</b>   <b>ipv4 labeled-unicast</b>   <b>ipv4 tunnel</b>   <b>ipv4 mdt</b>   <b>ipv6 unicast</b>   <b>ipv6 multicast</b>   <b>ipv6 labeled-unicast</b> <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast	Enters neighbor address family configuration mode for the specified address family.

	Command or Action	Purpose
Step 5	<p><code>soft-reconfiguration inbound [always]</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)#  soft-reconfiguration inbound always</p>	<p>Configures the software to store updates received from a specified neighbor. Soft reconfiguration inbound causes the software to store the original unmodified route in addition to a route that is modified or filtered. This allows a “soft clear” to be performed after the inbound policy is changed.</p> <p>Soft reconfiguration enables the software to store the incoming updates before apply policy if route refresh is not supported by the peer (otherwise a copy of the update is not stored). The <b>always</b> keyword forces the software to store a copy even when route refresh is supported by the peer.</p>
Step 6	<p><code>end</code>  OR  <code>commit</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end  OR  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Distributed BGP

Perform this task to configure distributed BGP. Configuring distributed BGP includes starting the speaker process and allocating the speaker process to a neighbor.

### Restrictions

If BGP is running in standalone mode, the **clear bgp current-mode** or **clear bgp vrf all \*** command must be used to switch from standalone mode to distributed mode.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **distributed speaker** *id*
4. **commit**

5. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **vpn4 unicast** | **vpn6 unicast** }
6. **exit**
7. **neighbor** *ip-address*
8. **remote-as** *as-number*
9. **speaker-id** *id*
10. **address-family** { **ipv4 unicast** | **ipv4 multicast** | **ipv4 labeled-unicast** | **ipv4 tunnel** | **ipv4 mdt** | **ipv6 unicast** | **ipv6 multicast** | **ipv6 labeled-unicast** | **vpn4 unicast** | **vpn6 unicast** }
11. **end**
12. **clear bgp current-mode**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>distributed speaker</b> <i>id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# distributed speaker 2	Specifies the speaker process to start.
Step 4	<b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# commit	Saves the configuration changes to the running configuration file and remains within the configuration session.
Step 5	<b>address-family</b> { <b>ipv4 unicast</b>   <b>ipv4 multicast</b>   <b>ipv4 tunnel</b>   <b>ipv4 mdt</b>   <b>ipv6 unicast</b>   <b>ipv6 multicast</b>   <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters neighbor address family configuration mode for the specified address family.
Step 6	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit	Exits address family mode.

	Command or Action	Purpose
Step 7	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp)# neighbor  172.168.40.24</p>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 8	<p><b>remote-as</b> <i>as-number</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as  2002</p>	Creates a neighbor and assigns a remote autonomous system number to it.
Step 9	<p><b>speaker-id</b> <i>id</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)#  speaker-id 2</p>	Allocates a neighbor to a specified speaker process.
Step 10	<p><b>address-family</b> {<i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 labeled-unicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>ipv6 labeled-unicast</i>   <i>vpn4 unicast</i>   <i>vpn6 unicast</i>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)#  address-family ipv4 unicast</p>	Enters neighbor address family configuration mode for the specified address family.
Step 11	<p><b>end</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul>
Step 12	<p><b>clear bgp current-mode</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router# clear bgp current-mode</p>	Switches from standalone mode to distributed mode.

## Configuring a VPN Routing and Forwarding Instance in BGP

The following tasks are used to configure a VPN routing and forwarding (VRF) instance in BGP:

- [Defining the Virtual Routing and Forwarding Tables in Provider Edge Routers](#), page RC-96 (required)
- [Configuring the Route Distinguisher](#), page RC-98 (required)
- [Configuring PE-PE or PE-RR Interior BGP sessions](#), page RC-100 (optional)
- [Configuring Route Reflector to Hold Routes That Have a Defined Set of RT Communities](#), page RC-103 (optional)
- [Configuring BGP as a PE-CE Protocol](#), page RC-105 (optional)
- [BGP Load Balancing](#), page RC-109 (optional)
- [Redistribution of IGP to BGP](#), page RC-111 (optional)

### Defining the Virtual Routing and Forwarding Tables in Provider Edge Routers

Perform this task to define the VPN routing and forwarding (VRF) tables in the provider edge (PE) routers.

#### SUMMARY STEPS

1. **configure**
2. **vrf** *vrf-name*
3. **address-family** {*ipv4* | *ipv6*} **unicast**
4. **maximum prefix** *maximum* [*threshold*]
5. **import route-policy** *policy-name*
6. **import route-target** [*as-number:nn* | *ip-address:nn*]
7. **export route-policy** *policy-name*
8. **export route-target** [*as-number:nn* | *ip-address:nn*]
9. **end**  
or  
**commit**

#### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# vrf vrf_pe	Configures a VRF instance.

	Command or Action	Purpose
Step 3	<p><b>address-family</b> {<i>ipv4</i>   <i>ipv6</i>} <i>unicast</i></p> <p><b>Example:</b>                      RP/0/RP0/CPU0:router(config-vrf)#                      address-family ipv4 unicast</p>	<p>Enters address family configuration mode for the specified address family.</p>
Step 4	<p><b>maximum prefix</b> <i>maximum</i> [<i>threshold</i>]</p> <p><b>Example:</b>                      RP/0/RP0/CPU0:router(config-vrf-af)# maximum                      prefix 2300</p>	<p>Configures a limit to the number of prefixes allowed in a VRF table.</p> <p>A maximum number of routes is applicable only to dynamic routing protocols and not to static or connected routes.</p> <p>You can specify a threshold percentage of the prefix limit using the <i>mid-threshold</i> argument.</p>
Step 5	<p><b>import route-policy</b> <i>policy-name</i></p> <p><b>Example:</b>                      RP/0/RP0/CPU0:router(config-vrf-af)# import                      route-policy <i>policy_a</i></p>	<p>Provides finer control over what gets imported into a VRF. This import filter discards prefixes that do not match the specified <i>policy-name</i> argument.</p>
Step 6	<p><b>import route-target</b> [<i>as-number:nn</i>   <i>ip-address:nn</i>]</p> <p><b>Example:</b>                      RP/0/RP0/CPU0:router(config-vrf-af)# import                      route-target 234:222</p>	<p>Specifies a list of route target (RT) extended communities. Only prefixes that are associated with the specified import route target extended communities are imported into the VRF.</p>
Step 7	<p><b>export route-policy</b> <i>policy-name</i></p> <p><b>Example:</b>                      RP/0/RP0/CPU0:router(config-vrf-af)# export                      route-policy <i>policy_b</i></p>	<p>Provides finer control over what gets exported into a VRF. This export filter discards prefixes that do not match the specified <i>policy-name</i> argument.</p>

	Command or Action	Purpose
Step 8	<p><b>export route-target</b> [<i>as-number:nn</i>   <i>ip-address:nn</i>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-vrf-af)# export route-target 123;234</p>	<p>Specifies a list of route target extended communities. Export route target communities are associated with prefixes when they are advertised to remote PEs. The remote PEs import them into VRFs which have import RTs that match these exported route target communities.</p>
Step 9	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-vrf-af)# end or RP/0/RP0/CPU0:router(config-vrf-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring the Route Distinguisher

The route distinguisher (RD) makes prefixes unique across multiple VPN routing and forwarding (VRF) instances.

Perform this task to configure the RD.

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- bgp router-id** *ip-address*
- vrf** *vrf-name*
- rd** {*as-number:nn* | *ip-address:nn* | **auto**}
- end**  
or  
**commit**



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode allowing you to configure the BGP routing process.
Step 3	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# bgp router-id 10.0.0.0	Configures a fixed router ID for the BGP-speaking router.
Step 4	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_pe	Configures a VRF instance.

Command or Action	Purpose
<p><b>Step 5</b></p> <pre>rd {as-number:nn   ip-address:nn   auto}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf)# rd  345:567</p>	<p>Configures the route distinguisher.</p> <p>Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</p> <p>Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation. The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</p>
<p><b>Step 6</b></p> <pre>end  OR  commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf)# end  OR  RP/0/RP0/CPU0:router(config-bgp-vrf)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before  exiting(yes/no/cancel)?  [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring PE-PE or PE-RR Interior BGP sessions

To enable BGP to carry VPN reachability information between provider edge (PE) routers you must configure the PE-PE interior BGP (IBGP) sessions. A PE uses VPN information carried from the remote PE router to determine VPN connectivity and the label value to be used so the remote (egress) router can demultiplex the packet to the correct VPN during packet forwarding.

The PE-PE, PE-route reflector (RR) IBGP sessions are defined to all PE and RR routers that participate in the VPNs configured in the PE router.

Perform this task to configure PE-PE IBGP sessions and to configure global VPN options on a PE.

### SUMMARY STEPS

1. **configure**
2. **router bgp as-number**

3. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
4. **bgp dampening** [*half-life* [*reuse suppress max-suppress-time*] | **route-policy** *route-policy-name*]
5. **bgp client-to-client reflection disable**
6. **exit**
7. **neighbor** *ip-address*
8. **remote-as** *as-number*
9. **description** *text*
10. **password** { **clear** | **encrypted** } *password*
11. **shutdown**
12. **timers keepalive** *hold-time*
13. **update-source** *interface-type interface-number*
14. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
15. **route-policy** *route-policy-name* **in**
16. **route-policy** *route-policy-name* **out**
17. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family vpn4 unicast	Enters VPN address family configuration mode.
Step 4	<b>bgp dampening</b> [ <i>half-life</i> [ <i>reuse suppress max-suppress-time</i> ]   <b>route-policy</b> <i>route-policy-name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# bgp dampening 15 1200 110 96	Enables BGP route dampening.

	Command or Action	Purpose
Step 5	<b>bgp client-to-client reflection disable</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection disable	Disables reflection of routes between route reflection clients using a BGP route reflector.
Step 6	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 7	<b>neighbor ip-address</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.16.1.1	Configures a PE IBGP neighbor.
Step 8	<b>remote-as as-number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1	Assigns the neighbor a remote autonomous system number.
Step 9	<b>description text</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# description neighbor 172.16.1.1	(Optional) Provides a description of the neighbor. The description is used to save comments and does not affect software function.
Step 10	<b>password {clear   encrypted} password</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# password encrypted 123abc	Enables Message Digest 5 (MD5) authentication on the TCP connection between the two BGP neighbors.
Step 11	<b>shutdown</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown	Terminates any active sessions for the specified neighbor and removes all associated routing information.
Step 12	<b>timers keepalive hold-time</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 12000 200	Set the timers for the BGP neighbor.
Step 13	<b>update-source interface-type interface-number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# update-source gigabitEthernet 0/1/5/0	Allows IBGP sessions to use the primary IP address from a specific interface as the local address when forming an IBGP session with a neighbor.

	Command or Action	Purpose
Step 14	<b>address-family</b> { <i>vpn4 unicast</i>   <i>vpn6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family vpn4 unicast	Enters VPN neighbor address family configuration mode.
Step 15	<b>route-policy</b> <i>route-policy-name in</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pe-pe-vpn-in in	Specifies a routing policy for an inbound route. The policy can be used to filter routes or modify route attributes.
Step 16	<b>route-policy</b> <i>route-policy-name out</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pe-pe-vpn-out out	Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.
Step 17	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end or RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Route Reflector to Hold Routes That Have a Defined Set of RT Communities

A provider edge (PE) needs to hold the routes that match the import route targets (RTs) of the VPNs configured on it. The PE router can discard all other VPNv4 (Cisco XR 12000 Series Router and Cisco CRS-1) and VPNv6 (Cisco XR 12000 Series Router only) routes. But, a route reflector (RR) must retain all VPNv4 and VPNv6 routes, because it might peer with PE routers and different PEs might require different RT-tagged VPNv4 and VPNv6 routes (making RRs non-scalable). You can configure an RR to only hold routes that have a defined set of RT communities. Also, a number of the RRs can be configured to service a different set of VPNs (thereby achieving some scalability). A PE is then made to peer with all RRs that service the VRFs configured on the PE. When a new VRF is configured with an RT for which the PE does not already hold routes, the PE issues route refreshes to the RRs and retrieves the relevant VPN routes.

**Note**

Note that this process can be more efficient if the PE-RR session supports extended community outbound route filter (ORF).

Perform this task to configure a reflector to retain routes tagged with specific RTs.

**SUMMARY STEPS**

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { *vpn4 unicast* | *vpn6 unicast* }
4. **retain route-target** { *all* | *route-policy route-policy-name* }
5. **end**  
or  
**commit**

**DETAILED STEPS**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <i>vpn4 unicast</i>   <i>vpn6 unicast</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family vpn4 unicast	Enters VPN address family configuration mode.

	Command or Action	Purpose
Step 4	<p><b>retain route-target</b> {<b>all</b>   <b>route-policy</b> <i>route-policy-name</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# retain route-target route-policy rr_ext-comm</p>	<p>Configures a reflector to retain routes tagged with particular RTs. Use the <i>route-policy-name</i> argument for the policy name that lists the extended communities that a path should have in order for the RR to retain that path.</p> <p><b>Note</b> The <b>all</b> keyword is not required, because this is the default behavior of a route reflector.</p>
Step 5	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP as a PE-CE Protocol

Perform this task to configure BGP on the PE and establish PE-CE communication using BGP.

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- vrf** *vrf-name*
- bgp router-id** *ip-address*
- label-allocation-mode per-ce**
- address-family** {**ipv4** | **ipv6**} **unicast**
- network** {*ip-address/prefix-length* | *ip-address mask*}
- aggregate-address** *address/mask-length*
- exit**
- neighbor** *ip-address*
- remote-as** *as-number*
- password** {**clear** | **encrypted**} *password*

13. **ebgp-multihop** [*ttl-value*]
14. **address-family** {*ipv4* {*unicast* | *labeled-unicast*} | *ipv6 unicast*}
15. **site-of-origin** [*as-number:nn* | *ip-address:nn*]
16. **as-override**
17. **allowas-in** [*as-occurrence-number*]
18. **route-policy** *route-policy-name* **in**
19. **route-policy** *route-policy-name* **out**
20. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_pe_2	Enables BGP routing for a particular VRF on the PE router.
Step 4	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf)# bgp router-id 172.16.9.9	Configures a fixed router ID for a BGP-speaking router.
Step 5	<b>label-allocation-mode</b> <b>per-ce</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce	Configures the per-CE label allocation mode to avoid an extra lookup on the PE router and conserve label space (per-prefix is the default label allocation mode). In this mode, the PE router allocates one label for every immediate next-hop (in most cases, this would be a CE router). This label is directly mapped to the next hop, so there is no VRF route lookup performed during data forwarding. However, the number of labels allocated would be one for each CE rather than one for each VRF. Because BGP knows all the next hops, it assigns a label for each next hop (not for each PE-CE interface). When the outgoing interface is a multiaccess interface and the media access control (MAC) address of the neighbor is not known, Address Resolution Protocol (ARP) is triggered during packet forwarding.



	Command or Action	Purpose
Step 6	<p><b>address-family</b> {<i>ipv4</i>   <i>ipv6</i>} <i>unicast</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf)#  address-family ipv4 unicast</p>	Enters VPN address family configuration mode.
Step 7	<p><b>network</b> {<i>ip-address/prefix-length</i>   <i>ip-address mask</i>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-af)#  network 172.16.5.5</p>	Originates a network prefix in the address family table in the VRF context.
Step 8	<p><b>aggregate-address</b> <i>address/mask-length</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-af)#  aggregate-address 10.0.0.0/24</p>	Configures aggregation in the VRF address family context to summarize routing information to reduce the state maintained in the core. This summarization introduces some inefficiency in the PE edge, because an additional lookup is required to determine the ultimate next hop for a packet. When configured, a summary prefix is advertised instead of a set of component prefixes, which are more specifics of the aggregate. The PE advertises only one label for the aggregate. Because component prefixes could have different next hops to CEs, an additional lookup has to be performed during data forwarding.
Step 9	<p><b>exit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-af)# exit</p>	Exits the current configuration mode.
Step 10	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor  10.0.0.0</p>	Configures a CE neighbor. The <i>ip-address</i> argument must be a private address.
Step 11	<p><b>remote-as</b> <i>as-number</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  remote-as 2</p>	Configures the remote AS for the CE neighbor.
Step 12	<p><b>password</b> {<i>clear</i>   <i>encrypted</i>} <i>password</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  password encrypted 234xyz</p>	Enable Message Digest 5 (MD5) authentication on a TCP connection between two BGP neighbors.
Step 13	<p><b>ebgp-multihop</b> [<i>ttl-value</i>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  ebgp-multihop 55</p>	Configures the CE neighbor to accept and attempt BGP connections to external peers residing on networks that are not directly connected.

	Command or Action	Purpose
Step 14	<p><b>address-family</b> {<i>ipv4</i> {<i>unicast</i>   <i>labeled-unicast</i>}   <i>ipv6 unicast</i>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  address-family ipv4 unicast</p>	Enables an address family for the CE peer.
Step 15	<p><b>site-of-origin</b> [<i>as-number:nn</i>   <i>ip-address:nn</i>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  site-of-origin 234:111</p>	Configures the site-of-origin (SoO) extended community. Routes that are learned from this CE neighbor are tagged with the SoO extended community before being advertised to the rest of the PEs. SoO is frequently used to detect loops when as-override is configured on the PE router. If the prefix is looped back to the same site, the PE detects this and does not send the update to the CE.
Step 16	<p><b>as-override</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  as-override</p>	<p>Configures AS override on the PE router. This causes the PE router to replace the CE's ASN with its own (PE) ASN.</p> <p><b>Note</b> This loss of information could lead to routing loops; to avoid loops caused by as-override, use it in conjunction with site-of-origin.</p>
Step 17	<p><b>allowas-in</b> [<i>as-occurrence-number</i>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  allowas-in 5</p>	<p>Allows an AS path with the PE autonomous system number (ASN) a specified number of times.</p> <p>Hub and spoke VPN networks need the looping back of routing information to the HUB PE through the HUB CE. When this happens, due to the presence of the PE ASN, the looped-back information is dropped by the HUB PE. To avoid this, use the <b>allowas-in</b> command to allow prefixes even if they have the PEs ASN up to the specified number of times.</p>
Step 18	<p><b>route-policy</b> <i>route-policy-name in</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  route-policy pe_ce_in_policy in</p>	Specifies a routing policy for an inbound route. The policy can be used to filter routes or modify route attributes.

	Command or Action	Purpose
Step 19	<p><b>route-policy</b> <i>route-policy-name</i> <b>out</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  route-policy pe_ce_out_policy out</p>	Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.
Step 20	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)# en  d  or  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)#  commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## BGP Load Balancing

Perform this task to enable BGP multipath load sharing for external BGP (EBGP), interior BGP (IBGP), and EIBGP and to enable BGP to carry link bandwidth attribute of the DMZ link.

When the PE router includes the link bandwidth extended community in its updates to the remote PE through the Multiprotocol Interior BGP (MP-IBGP) session (either IPv4 or VPNv4), the remote PE automatically does load balancing if the **maximum-paths** command is enabled.

### SUMMARY STEPS

- configure**
- router** *as-number*
- vrf** *vrf-name*
- address-family** {*ipv4* | *ipv6*} **unicast**
- maximum-paths** {*ebgp* | *ibgp* | *eibgp*} *maximum* [*unequal-cost*]
- exit**
- neighbor** *ip-address*
- dmz-link-bandwidth**

```

9. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_a	Enables BGP routing for a particular VRF on the PE router.
Step 4	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast	Enters VPN address family configuration mode.
Step 5	<b>maximum-paths</b> { <i>ebgp</i>   <i>ibgp</i>   <i>eibgp</i> } <i>maximum</i> [ <i>unequal-cost</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf-af)# maximum-paths ebgp 3	Configures the maximum number of parallel routes that BGP installs in the routing table. <ul style="list-style-type: none"> <li>• <b>ebgp</b> <i>maximum</i>: Consider only EBGp paths for multipath.</li> <li>• <b>ibgp</b> <i>maximum</i> [<i>unequal-cost</i>]: Consider load balancing between IBGP learned paths.</li> <li>• <b>eibgp</b>: Consider both EBGp and IBGP learned paths for load balancing. EIBGP load balancing always does unequal-cost load balancing.</li> </ul> <p>When EIBGP is applied, EBGp or IBGP load balancing cannot be configured. EBGp and IBGP load balancing can coexist.</p>
Step 6	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf-af)# exit	Exits the current configuration mode.
Step 7	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 10.0.0.0	Configures a CE neighbor. The <i>ip-address</i> argument must be a private address.

	Command or Action	Purpose
Step 8	<p><b>dmz-link-bandwidth</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  dmz-link-bandwidth</p>	<p>Originates a demilitarized-zone (DMZ) link-bandwidth extended community for the link to an EBGP neighbor.</p>
Step 9	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# end  or  RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)#  commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistribution of IGP to BGP

Perform this task to configure redistribution of a protocol into the VRF address family.

Even if Interior Gateway Protocols (IGPs) are used as the PE-CE protocol, the import logic happens through BGP. Therefore, all IGP routes have to be imported into the BGP VRF table.

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- vrf** *vrf-name*
- address-family** {*ipv4* | *ipv6*} **unicast**
- redistribute connected** [*metric metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute eigrp** *process-id* [**match** {*external* | *internal*}] [*metric metric-value*] [**route-policy** *route-policy-name*]  
or  
**redistribute isis** *process-id* [**level** {*1* | *1-inter-area* | *2*}] [*metric metric-value*] [**route-policy** *route-policy-name*]  
or

```
redistribute ospf process-id [match { external [1 | 2] | internal | nssa-external [1 | 2]}] [metric
metric-value] [route-policy route-policy-name]
```

or

```
redistribute ospfv3 process-id [match { external [1 | 2] | internal | nssa-external [1 | 2]}] [metric
metric-value] [route-policy route-policy-name]
```

or

```
redistribute rip [metric metric-value] [route-policy route-policy-name]
```

or

```
redistribute static [metric metric-value] [route-policy route-policy-name]
```

6. **end**
- or
- commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp as-number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_a	Enables BGP routing for a particular VRF on the PE router.
Step 4	<b>address-family {ipv4   ipv6} unicast</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast	Enters VPN address family configuration mode.

	Command or Action	Purpose
Step 5	<pre> <b>redistribute connected</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute eigrp</b> <i>process-id</i> [<b>match</b> {<b>external</b>   <b>internal</b>}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute isis</b> <i>process-id</i> [<b>level</b> {<b>1</b>   <b>1-inter-area</b>   <b>2</b>}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute ospf</b> <i>process-id</i> [<b>match</b> {<b>external</b> [<b>1</b>   <b>2</b>]   <b>internal</b>   <b>nssa-external</b> [<b>1</b>   <b>2</b>]}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute ospfv3</b> <i>process-id</i> [<b>match</b> {<b>external</b> [<b>1</b>   <b>2</b>]   <b>internal</b>   <b>nssa-external</b> [<b>1</b>   <b>2</b>]}] [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute rip</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>] OR <b>redistribute static</b> [<b>metric</b> <i>metric-value</i>] [<b>route-policy</b> <i>route-policy-name</i>]  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf-af)# redistribute eigrp 23 </pre>	<p>Configures redistribution of a protocol into the VRF address family context.</p> <p>The <b>redistribute</b> command is used if BGP is not used between the PE-CE routers. If BGP is used between PE-CE routers, the IGP that is used has to be redistributed into BGP to establish VPN connectivity with other PE sites. Redistribution is also required for inter-table import and export.</p>
Step 6	<pre> <b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-vrf-af)# <b>end</b> OR RP/0/RP0/CPU0:router(config-bgp-vrf-af)# <b>commit</b> </pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: </pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Keychains for BGP

Keychains provide secure authentication by supporting different MAC authentication algorithms and provide graceful key rollover. Perform this task to configure keychains for BGP. This task is optional.



### Note

If a keychain is configured for a neighbor group or a session group, a neighbor using the group inherits the keychain. Values of commands configured specifically for a neighbor override inherited values.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **keychain** *name*
6. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 4	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.



	Command or Action	Purpose
Step 5	<p><b>keychain</b> <i>name</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain  kych_a</p>	Configures keychain-based authentication.
Step 6	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-nbr)# end  or  RP/0/RP0/CPU0:router(config-bgp-nbr)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring an MDT Address Family Session in BGP

Perform this task to configure an IPv4 multicast distribution tree (MDT) subaddress family identifier (SAFI) session in BGP.



### Note

The MDT address-family session is configured on provider edge (PE) routers to establish VPN peering sessions with customer edge (CE) routers and to establish inter-AS multicast VPN peering sessions. The MDT address family must be configured on each participating PE router.

The MDT SAFI is a transitive, multicast-capable connector attribute that is defined as an IPv4 address family in BGP. The MDT SAFI is designed to support inter-autonomous system (inter-AS) VPN peering sessions.

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- address-family** {*ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast*}
- exit**

5. **address-family ipv4 mdt**
6. **exit**
7. **neighbor ip-address**
8. **remote-as as-number**
9. **address-family {ipv4 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6 multicast | ipv6 labeled-unicast | vpnv4 unicast | vpnv6 unicast}**
10. **route-policy route-policy-name {in | out}**
11. **exit**
12. **address-family ipv4 mdt**
13. **route-policy route-policy-name {in | out}**
14. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp as-number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family {ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv6 unicast   ipv6 multicast   vpnv4 unicast   vpnv6 unicast}</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast	Enters address family configuration mode for the specified address family.
Step 4	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 5	<b>address-family ipv4 mdt</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 mdt	Specifies the multicast distribution tree (MDT) address family.

	Command or Action	Purpose
Step 6	<p><b>exit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit</p>	Exits the current configuration mode.
Step 7	<p><b>neighbor ip-address</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24</p>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 8	<p><b>remote-as as-number</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002</p>	Creates a neighbor and assigns a remote autonomous system number to it.
Step 9	<p><b>address-family {ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpv4 unicast   vpv6 unicast}</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast</p>	Enters neighbor address family configuration mode for the specified address family.
Step 10	<p><b>route-policy route-policy-name {in   out}</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# route-policy pe_ce_out_policy out</p>	Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.
Step 11	<p><b>exit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-af)# exit</p>	Exits the current configuration mode.
Step 12	<p><b>address-family ipv4 mdt</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 mdt</p>	Specifies the multicast distribution tree (MDT) address family.

Command or Action	Purpose
<p><b>Step 13</b> <code>route-policy route-policy-name {in   out}</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)#  route-policy pe_ce_out_policy out</p>	<p>Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.</p>
<p><b>Step 14</b> <code>end</code>  or  <code>commit</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-bgp-af)# end  or  RP/0/RP0/CPU0:router(config-bgp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Disabling a BGP Neighbor

Perform this task to administratively shut down a neighbor session without removing the configuration.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **shutdown**
5. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 127	Enters BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 4	<b>shutdown</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown	Disables all active sessions for the specified neighbor.
Step 5	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# end OR RP/0/RP0/CPU0:router(config-bgp-nbr)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Resetting Neighbors Using BGP Inbound Soft Reset

Perform this task to trigger an inbound soft reset of the specified address families for the specified group or neighbors. The group is specified by the *\**, *ip-address*, *as-number*, or **external** keywords and arguments.

Resetting neighbors is useful if you change the inbound policy for the neighbors or any other configuration that affects the sending or receiving of routing updates. If an inbound soft reset is triggered, BGP sends a REFRESH request to the neighbor if the neighbor has advertised the ROUTE\_REFRESH capability. To determine whether the neighbor has advertised the ROUTE\_REFRESH capability, use the **show bgp neighbors** command.

### SUMMARY STEPS

1. **show bgp neighbors**
2. **clear bgp** {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} {\* | ip-address | as as-number | external} soft [in [prefix-filter] | out]

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show bgp neighbors</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp neighbors</p>	Verifies that received route refresh capability from the neighbor is enabled.
Step 2	<pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all   tunnel   mdt}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast   mdt   tunnel}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} {*   ip-address   as as-number   external} soft [in [prefix-filter]   out]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.1 soft in</p>	Soft resets a BGP neighbor. <ul style="list-style-type: none"> <li>• The <i>*</i> keyword resets all BGP neighbors.</li> <li>• The <i>ip-address</i> argument specifies the address of the neighbor to be reset.</li> <li>• The <i>as-number</i> argument specifies that all neighbors that match the autonomous system number be reset.</li> <li>• The <b>external</b> keyword specifies that all external neighbors are reset.</li> </ul>

## Resetting Neighbors Using BGP Outbound Soft Reset

Perform this task to trigger an outbound soft reset of the specified address families for the specified group or neighbors. The group is specified by the *\**, *ip-address*, *as-number*, or **external** keywords and arguments.

Resetting neighbors is useful if you change the outbound policy for the neighbors or any other configuration that affects the sending or receiving of routing updates.

If an outbound soft reset is triggered, BGP resends all routes for the address family to the given neighbors.

To determine whether the neighbor has advertised the ROUTE\_REFRESH capability, use the **show bgp neighbors** command.

## SUMMARY STEPS

1. **show bgp neighbors**
2. **clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} {\* | ip-address | as as-number | external} soft out**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show bgp neighbors</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp neighbors</p>	Verifies that received route refresh capability from the neighbor is enabled.
Step 2	<pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all   tunnel   mdt}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast   mdt   tunnel}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} {*   ip-address   as as-number   external} soft out</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.2 soft out</p>	Soft resets a BGP neighbor. <ul style="list-style-type: none"> <li>• The <b>*</b> keyword resets all BGP neighbors.</li> <li>• The <i>ip-address</i> argument specifies the address of the neighbor to be reset.</li> <li>• The <i>as-number</i> argument specifies that all neighbors that match the autonomous system number be reset.</li> <li>• The <b>external</b> keyword specifies that all external neighbors are reset.</li> </ul>

## Resetting Neighbors Using BGP Hard Reset

Perform this task to reset neighbors using a hard reset. A hard reset removes the TCP connection to the neighbor, removes all routes received from the neighbor from the BGP table, and then re-establishes the session with the neighbor. If the **graceful** keyword is specified, the routes from the neighbor are not removed from the BGP table immediately, but are marked as stale. After the session is re-established, any stale route that has not been received again from the neighbor is removed.

## SUMMARY STEPS

1. **clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast} {\* | ip-address | as as-number | external} [graceful] soft [in [prefix-filter] | out]**

## DETAILED STEPS

Command or Action	Purpose
<p><b>Step 1</b></p> <pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all   tunnel   mdt}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast   mdt   tunnel}   vpv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpv6 unicast} {*   ip-address   as as-number   external} [graceful] soft [in [prefix-filter]   out]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.3</p>	<p>Clears a BGP neighbor.</p> <ul style="list-style-type: none"> <li>The <code>*</code> keyword resets all BGP neighbors.</li> <li>The <code>ip-address</code> argument specifies the address of the neighbor to be reset.</li> <li>The <code>as-number</code> argument specifies that all neighbors that match the autonomous system number be reset.</li> <li>The <code>external</code> keyword specifies that all external neighbors are reset.</li> </ul> <p>The <code>graceful</code> keyword specifies a graceful restart.</p>

## Clearing Caches, Tables, and Databases

Perform this task to remove all contents of a particular cache, table, or database. The **clear bgp** command resets the sessions of the specified group of neighbors (hard reset); it removes the TCP connection to the neighbor, removes all routes received from the neighbor from the BGP table, and then re-establishes the session with the neighbor. Clearing a cache, table, or database can become necessary when the contents of the particular structure have become, or are suspected to be, invalid.

## SUMMARY STEPS

- clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpv6 unicast} ip-address**
- clear bgp external**
- clear bgp \***



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>clear bgp {ipv4 {unicast   multicast   labeled-unicast   all   tunnel   mdt}   ipv6 {unicast   multicast   all   labeled-unicast}   all {unicast   multicast   all   labeled-unicast   mdt   tunnel}   vpnv4 unicast   vrf {vrf-name   all} {ipv4 {unicast   labeled-unicast}   ipv6 unicast}   vpnv6 unicast} ip-address</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp ipv4 172.20.1.1</p>	Clears a specified neighbor.
Step 2	<pre>clear bgp external</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp external</p>	Clears all external peers.
Step 3	<pre>clear bgp *</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear bgp *</p>	Clears all BGP neighbors.

## Displaying System and Network Statistics

Perform this task to display specific statistics, such as the contents of BGP routing tables, caches, and databases. Information provided can be used to determine resource usage and solve network problems. You can also display information about node reachability and discover the routing path that the packets of your device are taking through the network.

## SUMMARY STEPS

1. `show bgp cidr-only`
2. `show bgp community community-list [exact-match]`
3. `show bgp regexp regular-expression`
4. `show bgp`
5. `show bgp neighbors ip-address [advertised-routes | dampened-routes | flap-statistics | performance-statistics | received prefix-filter | routes]`
6. `show bgp paths`
7. `show bgp neighbor-group group-name configuration`
8. `show bgp summary`

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show bgp cidr-only</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp cidr-only</p>	Displays routes with nonnatural network masks (classless interdomain routing [CIDR]) routes.
Step 2	<pre>show bgp community community-list [exact-match]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp community 1081:5 exact-match</p>	Displays routes that match the specified BGP community.
Step 3	<pre>show bgp regexp regular-expression</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp regexp "^3 "</p>	Displays routes that match the specified autonomous system path regular expression.
Step 4	<pre>show bgp</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp</p>	Displays entries in the BGP routing table.
Step 5	<pre>show bgp neighbors ip-address [advertised-routes   dampened-routes   flap-statistics   performance-statistics   received prefix-filter   routes]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.1</p>	<p>Displays information about the BGP connection to the specified neighbor.</p> <ul style="list-style-type: none"> <li>• The <b>advertised-routes</b> keyword displays all routes the router advertised to the neighbor.</li> <li>• The <b>dampened-routes</b> keyword displays the dampened routes that are learned from the neighbor.</li> <li>• The <b>flap-statistics</b> keyword displays flap statistics of the routes learned from the neighbor.</li> <li>• The <b>performance-statistics</b> keyword displays performance statistics relating to work done by the BGP process for this neighbor.</li> <li>• The <b>received prefix-filter</b> keyword and argument display the received prefix list filter.</li> <li>• The <b>routes</b> keyword displays routes learned from the neighbor.</li> </ul>
Step 6	<pre>show bgp paths</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp paths</p>	Displays all BGP paths in the database.

	Command or Action	Purpose
Step 7	<b>show bgp neighbor-group</b> <i>group-name</i> <b>configuration</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show bgp neighbor-group group_1 configuration	Displays the effective configuration for a specified neighbor group, including any configuration inherited by this neighbor group.
Step 8	<b>show bgp summary</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show bgp summary	Displays the status of all BGP connections.

## Displaying BGP Process Information

Perform this task to display specific BGP process information.

### SUMMARY STEPS

1. **show bgp process**
2. **show bgp ipv4 unicast summary**
3. **show bgp vpnv4 unicast summary**
4. **show bgp vrf {vrf-name | all}**
5. **show bgp process detail**
6. **show bgp summary**
7. **show placement program bgp**
8. **show placement program brib**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>show bgp process</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show bgp process	Displays status and summary information for the BGP process. The output shows various global and address family-specific BGP configurations. A summary of the number of neighbors, update messages, and notification messages sent and received by the process is also displayed.
Step 2	<b>show bgp ipv4 unicast summary</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show bgp ipv4 unicast summary	Displays a summary of the neighbors for the IPv4 unicast address family.
Step 3	<b>show bgp vpnv4 unicast summary</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show bgp vpnv4 unicast summary	Displays a summary of the neighbors for the VPNv4 unicast address family.

	Command or Action	Purpose
Step 4	<pre>show bgp vrf (vrf-name   all)</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp vrf vrf_A</p>	Displays BGP VPN virtual routing and forwarding (VRF) information.
Step 5	<pre>show bgp process detail</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp processes detail</p>	Displays detailed process information including the memory used by each of various internal structure types.
Step 6	<pre>show bgp summary</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp summary</p>	Displays the status of all BGP connections.
Step 7	<pre>show placement program bgp</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show placement program bgp</p>	Displays BGP program information. <ul style="list-style-type: none"> <li>• If a program is shown as having ‘rejected locations’ (for example, locations where program cannot be placed), the locations in question can be viewed using the <b>show placement program bgp</b> command.</li> <li>• If a program has been placed but not started, the amount of elapsed time since the program was placed is displayed in the Waiting to start column.</li> </ul>
Step 8	<pre>show placement program brib</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show placement program brib</p>	Displays bRIB program information. <ul style="list-style-type: none"> <li>• If a program is shown as having ‘rejected locations’ (for example, locations where program cannot be placed), the locations in question can be viewed using the <b>show placement program bgp</b> command.</li> <li>• If a program has been placed but not started, the amount of elapsed time since the program was placed is displayed in the Waiting to start column.</li> </ul>

## Monitoring BGP Update Groups

This task displays information related to the processing of BGP update groups.

### SUMMARY STEPS

1. **show bgp [ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} | vpnv6 unicast] update-group [neighbor ip-address | process-id.index [summary | performance-statistics]]**

## DETAILED STEPS

Command or Action	Purpose
<p><b>Step 1</b></p> <pre>show bgp [ipv4 {unicast   multicast   labeled-unicast   all   tunnel   mdt}   ipv6 {unicast   all   labeled-unicast}   all {unicast   multicast   all   mdt   labeled-unicast   tunnel}   vpnv4 unicast   vrf {vrf-name   all} [ipv4 {unicast   labeled-unicast}   ipv6 unicast]   vpnv6 unicast] update-group [neighbor ip-address   process-id.index [summary   performance-statistics]]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show bgp update-group 0.0</p>	<p>Displays information about BGP update groups.</p> <ul style="list-style-type: none"> <li>The <i>ip-address</i> argument displays the update groups to which that neighbor belongs.</li> <li>The <i>process-id.index</i> argument selects a particular update group to display and is specified as follows: process ID (dot) index. Process ID range is from 0 to 254. Index range is from 0 to 4294967295.</li> <li>The <b>summary</b> keyword displays summary information for neighbors in a particular update group.</li> <li>If no argument is specified, this command displays information for all update groups (for the specified address family).</li> <li>The <b>performance-statistics</b> keyword displays performance statistics for an update group.</li> </ul>

## Configuration Examples for Implementing BGP on Cisco IOS XR Software

This section provides the following configuration examples:

- [Enabling BGP: Example, page RC-127](#)
- [Displaying BGP Update Groups: Example, page RC-129](#)
- [BGP Neighbor Configuration: Example, page RC-129](#)
- [BGP Confederation: Example, page RC-130](#)
- [BGP Route Reflector: Example, page RC-131](#)

### Enabling BGP: Example

The following shows how to enable BGP.

```
prefix-set static
 2020::/64,
 2012::/64,
 10.10.0.0/16,
 10.2.0.0/24
end-set

route-policy pass-all
 pass
end-policy
route-policy set_next_hop_agg_v4
 set next-hop 10.0.0.1
end-policy
```

```

route-policy set_next_hop_static_v4
  if (destination in static) then
    set next-hop 10.1.0.1
  else
    drop
  endif
end-policy

route-policy set_next_hop_agg_v6
  set next-hop 2003::121
end-policy

route-policy set_next_hop_static_v6
  if (destination in static) then
    set next-hop 2011::121
  else
    drop
  endif
end-policy

router bgp 65000
  bgp fast-external-fallover disable
  bgp confederation peers
    65001
    65002
  bgp confederation identifier 1
  bgp router-id 1.1.1.1
  address-family ipv4 unicast
    aggregate-address 10.2.0.0/24 route-policy set_next_hop_agg_v4
    aggregate-address 10.3.0.0/24
    redistribute static route-policy set_next_hop_static_v4
  address-family ipv4 multicast
    aggregate-address 10.2.0.0/24 route-policy set_next_hop_agg_v4
    aggregate-address 10.3.0.0/24
    redistribute static route-policy set_next_hop_static_v4
  address-family ipv6 unicast
    aggregate-address 2012::/64 route-policy set_next_hop_agg_v6
    aggregate-address 2013::/64
    redistribute static route-policy set_next_hop_static_v6
  address-family ipv6 multicast
    aggregate-address 2012::/64 route-policy set_next_hop_agg_v6
    aggregate-address 2013::/64
    redistribute static route-policy set_next_hop_static_v6
  neighbor 10.0.101.60
    remote-as 65000
    address-family ipv4 unicast
    address-family ipv4 multicast
  neighbor 10.0.101.61
    remote-as 65000
    address-family ipv4 unicast
    address-family ipv4 multicast
  neighbor 10.0.101.62
    remote-as 3
    address-family ipv4 unicast
    route-policy pass-all in
    route-policy pass-all out
    address-family ipv4 multicast
    route-policy pass-all in
    route-policy pass-all out
  neighbor 10.0.101.64
    remote-as 5
    update-source Loopback0
    address-family ipv4 unicast
    route-policy pass-all in

```

```
route-policy pass-all out
address-family ipv4 multicast
route-policy pass-all in
route-policy pass-all out
```

## Displaying BGP Update Groups: Example

The following is sample output from the **show bgp update-group** command run in EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp update-group

Update group for IPv4 Unicast, index 0.1:
  Attributes:
    Outbound Route map:rm
    Minimum advertisement interval:30
    Messages formatted:2, replicated:2
  Neighbors in this update group:
    10.0.101.92

Update group for IPv4 Unicast, index 0.2:
  Attributes:
    Minimum advertisement interval:30
    Messages formatted:2, replicated:2
  Neighbors in this update group:
    10.0.101.91
```

## BGP Neighbor Configuration: Example

The following example shows how BGP neighbors on an autonomous system are configured to share information. In the example, a BGP router is assigned to autonomous system 109, and two networks are listed as originating in the autonomous system. Then the addresses of three remote routers (and their autonomous systems) are listed. The router being configured shares information about networks 131.108.0.0 and 192.31.7.0 with the neighbor routers. The first router listed is in a different autonomous system; the second **neighbor** and **remote-as** commands specify an internal neighbor (with the same autonomous system number) at address 131.108.234.2; and the third **neighbor** and **remote-as** commands specify a neighbor on a different autonomous system.

```
route-policy pass-all
  pass
end-policy
router bgp 109
  address-family ipv4 unicast
    network 131.108.0.0 255.0.0.0
    network 192.31.7.0 255.0.0.0
    neighbor 131.108.200.1
      remote-as 167
    exit
  address-family ipv4 unicast
    route-policy pass-all in
    route-policy pass-out out
    neighbor 131.108.234.2
      remote-as 109
    exit
  address-family ipv4 unicast
    neighbor 150.136.64.19
      remote-as 99
    exit
  address-family ipv4 unicast
```

```
route-policy pass-all in
route-policy pass-all out
```

## BGP Confederation: Example

The following is a sample configuration that shows several peers in a confederation. The confederation consists of three internal autonomous systems with autonomous system numbers 6001, 6002, and 6003. To the BGP speakers outside the confederation, the confederation looks like a normal autonomous system with autonomous system number 666 (specified using the **bgp confederation identifier** command).

In a BGP speaker in autonomous system 6001, the **bgp confederation peers** command marks the peers from autonomous systems 6002 and 6003 as special eBGP peers. Hence, peers 171.69.232.55 and 171.69.232.56 get the local preference, next hop, and MED unmodified in the updates. The router at 160.69.69.1 is a normal eBGP speaker, and the updates received by it from this peer are just like a normal eBGP update from a peer in autonomous system 666.

```
router bgp 6001
  bgp confederation identifier 666
  bgp confederation peers
    6002
    6003
  exit
  address-family ipv4 unicast
    neighbor 171.69.232.55
    remote-as 6002
  exit
  address-family ipv4 unicast
    neighbor 171.69.232.56
    remote-as 6003
  exit
  address-family ipv4 unicast
    neighbor 160.69.69.1
    remote-as 777
```

In a BGP speaker in autonomous system 6002, the peers from autonomous systems 6001 and 6003 are configured as special eBGP peers. Peer 170.70.70.1 is a normal iBGP peer, and peer 199.99.99.2 is a normal eBGP peer from autonomous system 700.

```
router bgp 6002
  bgp confederation identifier 666
  bgp confederation peers
    6001
    6003
  exit
  address-family ipv4 unicast
    neighbor 170.70.70.1
    remote-as 6002
  exit
  address-family ipv4 unicast
    neighbor 171.69.232.57
    remote-as 6001
  exit
  address-family ipv4 unicast
    neighbor 171.69.232.56
    remote-as 6003
  exit
  address-family ipv4 unicast
    neighbor 199.99.99.2
    remote-as 700
```



```
exit
address-family ipv4 unicast
route-policy pass-all in
route-policy pass-all out
```

In a BGP speaker in autonomous system 6003, the peers from autonomous systems 6001 and 6002 are configured as special eBGP peers. Peer 200.200.200.200 is a normal eBGP peer from autonomous system 701.

```
router bgp 6003
  bgp confederation identifier 666
  bgp confederation peers
    6001
    6002
  exit
  address-family ipv4 unicast
  neighbor 171.69.232.57
  remote-as 6001
  exit
  address-family ipv4 unicast
  neighbor 171.69.232.55
  remote-as 6002
  exit
  address-family ipv4 unicast
  neighbor 200.200.200.200
  remote-as 701
  exit
  address-family ipv4 unicast
  route-policy pass-all in
  route-policy pass-all out
```

The following is a part of the configuration from the BGP speaker 200.200.200.205 from autonomous system 701 in the same example. Neighbor 171.69.232.56 is configured as a normal eBGP speaker from autonomous system 666. The internal division of the autonomous system into multiple autonomous systems is not known to the peers external to the confederation.

```
router bgp 701
  address-family ipv4 unicast
  neighbor 171.69.232.56
  remote-as 666
  exit
  address-family ipv4 unicast
  route-policy pass-all in
  route-policy pass-all out
  exit
  address-family ipv4 unicast
  neighbor 200.200.200.205
  remote-as 701
```

## BGP Route Reflector: Example

The following example shows how to use an address family to configure internal BGP peer 10.1.1.1 as a route reflector client for both unicast and multicast prefixes:

```
router bgp 140
  address-family ipv4 unicast
  neighbor 10.1.1.1
  remote-as 140
  address-family ipv4 unicast
  route-reflector-client
  exit
```

```
address-family ipv4 multicast
route-reflector-client
```

## BGP MDT Address Family Configuration: Example

The following example shows how to configure an MDT address family in BGP:

```
router bgp 10
  bgp router-id 10.0.0.2
  address-family vpnv4 unicast
  !
  address-family ipv4 mdt
  !
  neighbor 1.1.1.1
    remote-as 11
    address-family vpnv4 unicast
      route-policy pass-all in
      route-policy pass-all out
    !
    address-family ipv4 mdt
      route-policy pass-all in
      route-policy pass-all out
    !
  !
  !
```

## Where to Go Next

For detailed information about BGP commands, see *Cisco IOS XR Routing Command Reference* document.

# Additional References

The following sections provide references related to implementing BGP for Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
BGP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference</i>
Cisco Express Forwarding (CEF) commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR IP Addresses and Services Command Reference</i>
MPLS VPN configuration information.	<i>Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>
Bidirectional Forwarding Detection (BFD)	<i>Cisco IOS XR Interface and Hardware Configuration Guide</i> and <i>Cisco IOS XR Interface and Hardware Command Reference</i>
Task ID information.	<i>Configuring AAA Services on Cisco IOS XR Software</i> module of <i>Cisco IOS XR System Security Configuration Guide</i>

## Standards

Standards	Title
draft-bonica-tcp-auth-05.txt	<i>Authentication for TCP-based Routing and Management Protocols</i> , by R. Bonica, B. Weis, S. Viswanathan, A. Lange, O. Wheeler
draft-ietf-idr-bgp4-26.txt	<i>A Border Gateway Protocol 4</i> , by Y. Rekhter, T.Li, S. Hares
draft-ietf-idr-bgp4-mib-15.txt	<i>Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4)</i> , by J. Hass and S. Hares
draft-ietf-idr-cease-subcode-05.txt	<i>Subcodes for BGP Cease Notification Message</i> , by Enke Chen, V. Gillet
draft-ietf-idr-avoid-transition-00.txt	<i>Avoid BGP Best Path Transitions from One External to Another</i> , by Enke Chen, Srihari Sangli
draft-ietf-idr-as4bytes-12.txt	<i>BGP Support for Four-octet AS Number Space</i> , by Quaizar Vohra, Enke Chen
draft-nalawade-idr-mdt-safi-03.txt	<i>MDT SAFI</i> , by Gargi Nalawade and Arjun Sreekantiah

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
RFC 1700	<i>Assigned Numbers</i>
RFC 1997	<i>BGP Communities Attribute</i>
RFC 2385	<i>Protection of BGP Sessions via the TCP MD5 Signature Option</i>
RFC 2439	<i>BGP Route Flap Damping</i>
RFC 2545	<i>Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing</i>
RFC 2796	<i>BGP Route Reflection - An Alternative to Full Mesh IBGP</i>
RFC 2858	<i>Multiprotocol Extensions for BGP-4</i>
RFC 2918	<i>Route Refresh Capability for BGP-4</i>
RFC 3065	<i>Autonomous System Confederations for BGP</i>
RFC 3392	<i>Capabilities Advertisement with BGP-4</i>
RFC 4271	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 4364	<i>BGP/MPLS IP Virtual Private Networks (VPNs)</i>
RFC 4724	<i>Graceful Restart Mechanism for BGP</i>

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>



# Implementing EIGRP on Cisco IOS XR Software

The Enhanced Interior Gateway Routing Protocol (EIGRP) is an enhanced version of IGRP developed by Cisco. EIGRP uses distance vector routing technology, which specifies that a router need not know all the router and link relationships for the entire network. Each router advertises destinations with a corresponding distance and upon receiving routes, adjusts the distance and propagates the information to neighboring routes.

This module describes the concepts and tasks you need to implement basic EIGRP configuration using Cisco IOS XR software.

For EIGRP configuration information related to the following features, see the “[Related Documents](#)” section of this module.

- Multiprotocol Label Switching (MPLS) Layer 3 Virtual Private Network (VPN)
- Site of Origin (SoO) Support



## Note

For more information about EIGRP on the Cisco IOS XR software and complete descriptions of the EIGRP commands listed in this module, see the “[Related Documents](#)” section of this module. To locate documentation for other commands that might appear while executing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing EIGRP on Cisco IOS XR Software

Release	Modification
Release 3.3.0	This feature was introduced on the Cisco CRS-1 and Cisco XR 12000 Series Router.
Release 3.4.0	Four-byte autonomous system (AS) number support was added.
Release 3.5.0	IPv6 address family and IPv6 VPN routing and forwarding (VRF) address family support was added.  IPv6 Provider Edge and IPv6 VPN Provider Edge Transport over Multiprotocol Label Switching Infrastructure support was added.
Release 3.6.0	No modification.

## Contents

- [Prerequisites for Implementing EIGRP on Cisco IOS XR Software](#), page RC-136
- [Restrictions for Implementing EIGRP on Cisco IOS XR Software](#), page RC-136

- [Information About Implementing EIGRP on Cisco IOS XR Software](#), page RC-136
- [How to Implement EIGRP on Cisco IOS XR Software](#), page RC-148
- [Configuration Examples for Implementing EIGRP on Cisco IOS XR Software](#), page RC-166
- [Additional References](#), page RC-168

## Prerequisites for Implementing EIGRP on Cisco IOS XR Software

The following are prerequisites for implementing EIGRP on Cisco IOS XR software:

- You must be in a user group associated with a task group that includes the proper task IDs for EIGRP commands. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of the *Cisco IOS XR System Security Configuration Guide*.

## Restrictions for Implementing EIGRP on Cisco IOS XR Software

The following restrictions are employed when running EIGRP on this version of Cisco IOS XR software:

- Only one instance of an EIGRP process is supported.
- Bidirectional Forwarding Detection (BFD) feature and the Simple Network Management Protocol (SNMP) MIB are not supported.
- Interface static routes are not automatically redistributed into EIGRP, because there are no network commands.
- Metric configuration (either through the **default-metric** command or a route policy) is required for redistribution of connected and static routes.
- Auto summary is disabled by default.
- Stub leak maps are not supported.
- Authentication is not supported.

## Information About Implementing EIGRP on Cisco IOS XR Software

To implement EIGRP, you need to understand the following concepts:

- [EIGRP Functional Overview](#), page RC-137
- [EIGRP Features](#), page RC-137
- [EIGRP Components](#), page RC-138
- [EIGRP Configuration Grouping](#), page RC-139
- [EIGRP Configuration Modes](#), page RC-139
- [Metric Weights for EIGRP Routing](#), page RC-141
- [Percentage of Link Bandwidth Used for EIGRP Packets](#), page RC-142

- [Floating Summary Routes for an EIGRP Process, page RC-142](#)
- [Split Horizon for an EIGRP Process, page RC-144](#)
- [Adjustment of Hello Interval and Hold Time for an EIGRP Process, page RC-145](#)
- [Stub Routing for an EIGRP Process, page RC-145](#)
- [Route Policy Options for an EIGRP Process, page RC-146](#)
- [EIGRP Layer 3 VPN PE-CE Site-of-Origin, page RC-147](#)
- [IPv6 and IPv6 VPN Provider Edge Support over MPLS and IP, page RC-148](#)

## EIGRP Functional Overview

Enhanced Interior Gateway Routing Protocol (EIGRP) is an interior gateway protocol suited for many different topologies and media. EIGRP scales well and provides extremely quick convergence times with minimal network traffic.

EIGRP has very low usage of network resources during normal operation. Only hello packets are transmitted on a stable network. When a change in topology occurs, only the routing table changes are propagated and not the entire routing table. Propagation reduces the amount of load the routing protocol itself places on the network. EIGRP also provides rapid convergence times for changes in the network topology.

The distance information in EIGRP is represented as a composite of available bandwidth, delay, load utilization, and link reliability with improved convergence properties and operating efficiency. The fine-tuning of link characteristics achieves optimal paths.

The convergence technology that EIGRP uses is based on research conducted at SRI International and employs an algorithm referred to as the Diffusing Update Algorithm (DUAL). This algorithm guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations. The convergence time with DUAL rivals that of any other existing routing protocol.

## EIGRP Features

EIGRP offers the following features:

- **Fast convergence**—The DUAL algorithm allows routing information to converge as quickly as any currently available routing protocol.
- **Partial updates**—EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature minimizes the bandwidth required for EIGRP packets.
- **Neighbor discovery mechanism**—This is a simple hello mechanism used to learn about neighboring routers. It is protocol independent.
- **Variable-length subnet masks (VLSMs).**
- **Arbitrary route summarization.**
- **Scaling**—EIGRP scales to large networks.

The following lists key features supported in the Cisco IOS XR implementation:

- Support for IPv4 and IPv6 address families.

- Provider Edge (PE)-Customer Edge (CE) protocol support with Site of Origin (SoO) and Border Gateway Protocol (BGP) cost community support.
- PECE protocol support for MPLS and L2TPv3-based-IP L3VPNs.

## EIGRP Components

EIGRP has the following four basic components:

- Neighbor discovery of neighbor recovery
- Reliable transport protocol
- DUAL finite state machine
- Protocol-dependent modules

Neighbor discovery or neighbor recovery is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery or neighbor recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, the Cisco IOS XR software can determine that a neighbor is alive and functioning. After this status is determined, the neighboring routers can exchange routing information.

The reliable transport protocol is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be sent reliably and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities (such as Ethernet) it is not necessary to send hello packets reliably to all neighbors individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets (such as updates) require acknowledgment, which is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when unacknowledged packets are pending. This provision helps to ensure that convergence time remains low in the presence of various speed links.

The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information (known as a metric) to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur. This is the process whereby a new successor is determined. The amount of time required to recompute the route affects the convergence time. Recomputation is processor intensive; it is advantageous to avoid unneeded recomputation. When a topology change occurs, DUAL tests for feasible successors. If there are feasible successors, it uses any it finds to avoid unnecessary recomputation.

The protocol-dependent modules are responsible for network layer protocol-specific tasks. An example is the EIGRP module, which is responsible for sending and receiving EIGRP packets that are encapsulated in IP. It is also responsible for parsing EIGRP packets and informing DUAL of the new information received. EIGRP asks DUAL to make routing decisions, but the results are stored in the IP routing table. EIGRP is also responsible for redistributing routes learned by other IP routing protocols.



## EIGRP Configuration Grouping

Cisco IOS XR software groups all EIGRP configuration under router EIGRP configuration mode, including interface configuration portions associated with EIGRP. To display EIGRP configuration in its entirety, use the **show running-config router eigrp** command. The command output displays the running configuration for the configured EIGRP instance, including the interface assignments and interface attributes.

## EIGRP Configuration Modes

The following examples show how to enter each of the configuration modes. From a mode, you can enter the ? command to display the commands available in that mode.

### Router Configuration Mode

The following example shows how to enter router configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)#
```

### VRF Configuration Mode

The following example shows how to enter VRF configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# vrf customer1
RP/0/RP0/CPU0:router(config-eigrp-vrf)#
```

### IPv4 Address Family Configuration Mode

The following example shows how to enter IPv4 address family configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)#
```

### IPv6 Address Family Configuration Mode

The following example shows how to enter IPv6 address family configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv6
RP/0/RP0/CPU0:router(config-eigrp-af)#
```

### IPv4 VRF Address Family Configuration Mode

The following example shows how to enter IPv4 VRF address family configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# vrf customer1
RP/0/RP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-vrf-af)#
```

### IPv6 VRF Address Family Configuration Mode

The following example shows how to enter IPv6 VRF address family configuration mode:

```
RP/0/RP0/CPU0:router# configuration
```

```
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# vrf customer1
RP/0/RP0/CPU0:router(config-eigrp-vrf)# address-family ipv6
RP/0/RP0/CPU0:router(config-eigrp-vrf-af)#
```

### Interface Configuration Mode

The following example shows how to enter interface configuration mode in IPv4 address family configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface POS0/3/0/0
RP/0/RP0/CPU0:router(config-eigrp-af-if)#
```

The following example shows how to enter interface configuration mode in IPv6 VRF configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# vrf customer1
RP/0/RP0/CPU0:router(config-eigrp-vrf)# address-family ipv6
RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# interface POS0/5/0/0
RP/0/RP0/CPU0:router(config-eigrp-vrf-af-if)#
```

## EIGRP Interfaces

EIGRP interfaces can be configured as either of the following types:

- **Active**—Advertises connected prefixes and forms adjacencies. This is the default type for interfaces.
- **Passive**—Advertises connected prefixes but does not form adjacencies. The **passive** command is used to configure interfaces as passive. Passive interfaces should be used sparingly for important prefixes, such as loopback addresses, that need to be injected into the EIGRP domain. If many connected prefixes need to be advertised, then the redistribution of connected routes with the appropriate policy should be used instead.

## Redistribution for an EIGRP Process

Routes from other protocols can be redistributed into EIGRP. A route policy can be configured along with the **redistribute** command. A metric is required, configured either through the **default-metric** command or under the route policy configured with the **redistribute** command to import routes into EIGRP.

A route policy allows the filtering of routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. When redistribution is configured under a VRF, EIGRP retrieves extended communities attached to the route in the routing information base (RIB). The SoO is used to filter out routing loops in the presence of MPLS VPN backdoor links.

## Metric Weights for EIGRP Routing

EIGRP uses the minimum bandwidth on the path to a destination network and the total delay to compute routing metrics. You can use the **metric weights** command to adjust the default behavior of EIGRP routing and metric computations. For example, this adjustment allows you to tune system behavior to allow for satellite transmission. EIGRP metric defaults have been carefully selected to provide optimal performance in most networks.

By default, the EIGRP composite metric is a 32-bit quantity that is a sum of the segment delays and lowest segment bandwidth (scaled and inverted) for a given route. For a network of homogeneous media, this metric reduces to a hop count. For a network of mixed media (FDDI, Ethernet, and serial lines running from 9600 bits per second to T1 rates), the route with the lowest metric reflects the most desirable path to a destination.

## Mismatched K Values

Mismatched K values (EIGRP metrics) can prevent neighbor relationships from being established and can negatively impact network convergence. The following example explains this behavior between two EIGRP peers (ROUTER-A and ROUTER-B).

The following error message is displayed in the console of ROUTER-B because the K values are mismatched:

```
RP/0/RP0/CPU0:Mar 13 08:19:55:eigrp[163]:%ROUTING-EIGRP-5-NBRCHANGE:IP-EIGRP(0) 1:Neighbor
11.0.0.20 (GigabitEthernet0/6/0/0) is down: K-value mismatch
```

Two scenarios occur in which this error message can be displayed:

- The two routers are connected on the same link and configured to establish a neighbor relationship. However, each router is configured with different K values.

The following configuration is applied to ROUTER-A. The K values are changed with the **metric weights** command. A value of 2 is entered for the *k1* argument to adjust the bandwidth calculation. The value of 1 is entered for the *k3* argument to adjust the delay calculation.

```
hostname ROUTER-A!
interface GigabitEthernet0/6/0/0
  ipv4 address 10.1.1.1 255.255.255.0

router eigrp 100
  metric weights 0 2 0 1 0 0
  interface GigabitEthernet0/6/0/0
```

The following configuration is applied to ROUTER-B. However, the **metric weights** command is not applied and the default K values are used. The default K values are 1, 0, 1, 0, and 0.

```
hostname ROUTER-B!
interface GigabitEthernet0/6/0/1
  ipv4 address 10.1.1.2 255.255.255.0

router eigrp 100
  interface GigabitEthernet0/6/0/1
```

The bandwidth calculation is set to 2 on ROUTER-A and set to 1 (by default) on ROUTER-B. This configuration prevents these peers from forming a neighbor relationship.

- The K-value mismatch error message can also be displayed if one of the two peers has transmitted a “goodbye” message and the receiving router does not support this message. In this case, the receiving router interprets this message as a K-value mismatch.

## The Goodbye Message

The goodbye message is a feature designed to improve EIGRP network convergence. The goodbye message is broadcast when an EIGRP routing process is shut down to inform adjacent peers about the impending topology change. This feature allows supporting EIGRP peers to synchronize and recalculate neighbor relationships more efficiently than would occur if the peers discovered the topology change after the hold timer expired.

The following message is displayed by routers that run a supported release when a goodbye message is received:

```
RP/0/RP0/CPU0:Mar 13 09:13:17:eigrp[163]:%ROUTING-EIGRP-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 10.0.0.20 (GigabitEthernet0/6/0/0) is down: Interface Goodbye received
```

A Cisco router that runs a software release that does not support the goodbye message can misinterpret the message as a K-value mismatch and display the following message:

```
RP/0/RP0/CPU0:Mar 13 09:13:17:eigrp[163]:%ROUTING-EIGRP-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 10.0.0.20 (GigabitEthernet0/6/0/0) is down: K-value mismatch
```



### Note

The receipt of a goodbye message by a non supporting peer does not disrupt normal network operation. The non supporting peer terminates the session when the hold timer expires. The sending and receiving routers reconverge normally after the sender reloads.

## Percentage of Link Bandwidth Used for EIGRP Packets

By default, EIGRP packets consume a maximum of 50 percent of the link bandwidth, as configured with the **bandwidth** interface configuration command. You might want to change that value if a different level of link utilization is required or if the configured bandwidth does not match the actual link bandwidth (it may have been configured to influence route metric calculations).

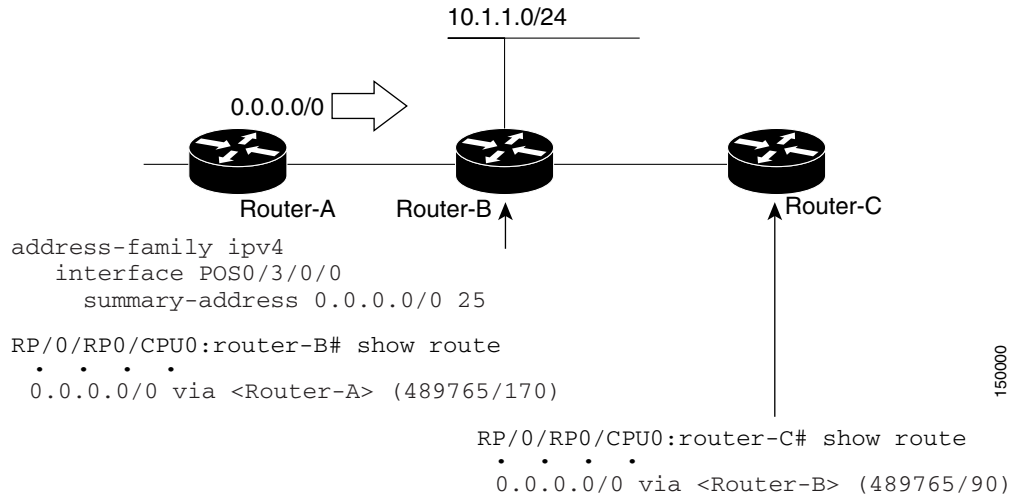
## Floating Summary Routes for an EIGRP Process

You can also use a floating summary route when configuring the **summary-address** command. The floating summary route is created by applying a default route and administrative distance at the interface level. The following scenario illustrates the behavior of this enhancement.

[Figure 1](#) shows a network with three routers, Router-A, Router-B, and Router-C. Router-A learns a default route from elsewhere in the network and then advertises this route to Router-B. Router-B is configured so that only a default summary route is advertised to Router-C. The default summary route is applied to interface 0/1 on Router-B with the following configuration:

```
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface POS0/3/0/0
RP/0/RP0/CPU0:router(config-eigrp-af-if)# summary-address 100.0.0.0 0.0.0.0
```

**Figure 1 Floating Summary Route Is Applied to Router-B**



The configuration of the default summary route on Router-B sends a 0.0.0.0/0 summary route to Router-C and blocks all other routes, including the 10.1.1.0/24 route, from being advertised to Router-C. However, this configuration also generates a local discard route on Router-B, a route for 0.0.0.0/0 to the null 0 interface with an administrative distance of 5. When this route is created, it overrides the EIGRP learned default route. Router-B is no longer able to reach destinations that it would normally reach through the 0.0.0.0/0 route.

This problem is resolved by applying a floating summary route to the interface on Router-B that connects to Router-C. The floating summary route is applied by relating an administrative distance to the default summary route on the interface of Router-B with the following statement:

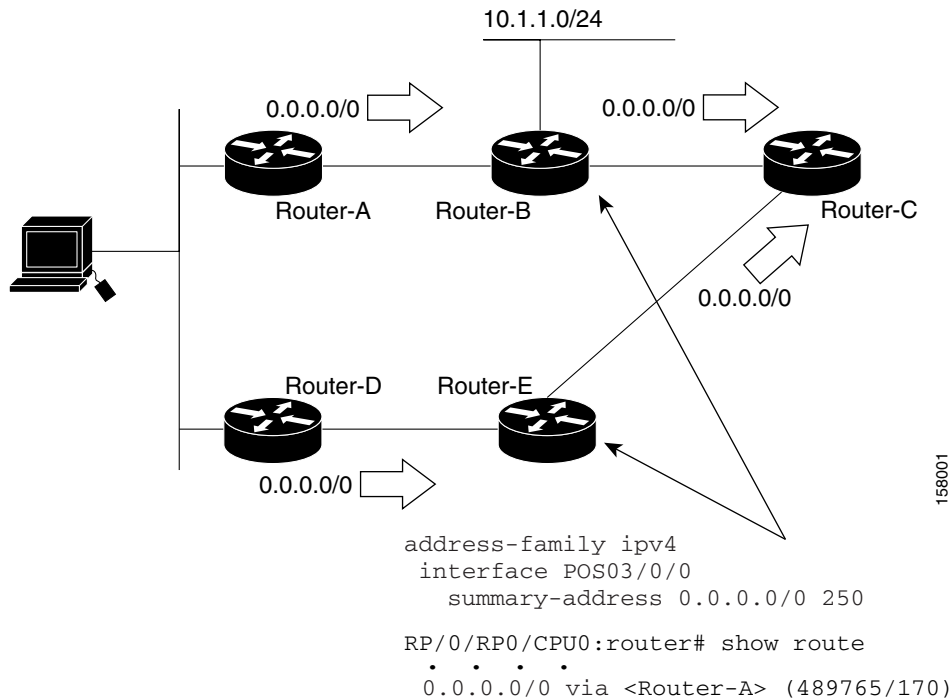
```
RP/0/RP0/CPU0:router(config-if)# summary-address 100 0.0.0.0 0.0.0.0 250
```

The administrative distance of 250, applied in the above statement, is now assigned to the discard route generated on Router-B. The 0.0.0.0/0, from Router-A, is learned through EIGRP and installed in the local routing table. Routing to Router-C is restored.

If Router-A loses the connection to Router-B, Router-B continues to advertise a default route to Router-C, which allows traffic to continue to reach destinations attached to Router-B. However, traffic destined for networks to Router-A or behind Router-A is dropped when the traffic reaches Router-B.

Figure 2 shows a network with two connections from the core: Router-A and Router-D. Both routers have floating summary routes configured on the interfaces connected to Router-C. If the connection between Router-E and Router-C fails, the network continues to operate normally. All traffic flows from Router-C through Router-B to the hosts attached to Router-A and Router-D.

**Figure 2 Floating Summary Route Applied for Dual-Homed Remotes**



However, if the link between Router-D and Router-E fails, the network may dump traffic into a black hole because Router-E continues to advertise the default route (0.0.0.0/0) to Router-C, as long as at least one link (other than the link to Router-C) to Router-E is still active. In this scenario, Router-C still forwards traffic to Router-E, but Router-E drops the traffic creating the black hole. To avoid this problem, you should configure the summary address with an administrative distance on only single-homed remote routers or areas in which only one exit point exists between the segments of the network. If two or more exit points exist (from one segment of the network to another), configuring the floating default route can cause a black hole to form.

## Split Horizon for an EIGRP Process

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, update and query packets are not sent for destinations for which this interface is the next hop. Controlling update and query packets in this manner reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router on any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks (such as Frame Relay and SMDS), situations can arise for which this behavior is less than ideal. For these situations, including networks in which you have EIGRP configured, you may want to disable split horizon.

## Adjustment of Hello Interval and Hold Time for an EIGRP Process

You can adjust the interval between hello packets and the hold time.

Routing devices periodically send hello packets to each other to dynamically learn of other routers on their directly attached networks. This information is used to discover neighbors and learn when neighbors become unreachable or inoperative. By default, hello packets are sent every 5 seconds.

You can configure the hold time on a specified interface for a particular EIGRP routing process designated by the autonomous system number. The hold time is advertised in hello packets and indicates to neighbors the length of time they should consider the sender valid. The default hold time is three times the hello interval, or 15 seconds.

## Stub Routing for an EIGRP Process

The EIGRP Stub Routing feature improves network stability, reduces resource usage, and simplifies stub router configuration.


Stub routing is commonly used in a hub-and-spoke network topology. In a hub-and-spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies in which the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router is connected to 100 or more remote routers. In a hub-and-spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP and configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The stub router responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” A router that is configured as a stub sends a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet informing it of the stub status does not query the stub router for any routes, and a router that has a stub peer does not query that peer. The stub router depends on the distribution router to send the proper updates to all peers.

[Figure 3](#) shows a simple hub-and-spoke configuration.

**Figure 3** Simple Hub-and-Spoke Network



The stub routing feature by itself does not prevent routes from being advertised to the remote router. In the example in [Figure 3](#), the remote router can access the corporate network and the Internet through the distribution router only. Having a full route table on the remote router, in this example, would serve no functional purpose because the path to the corporate network and the Internet would always be through the distribution router. The larger route table would only reduce the amount of memory required by the remote router. Bandwidth and memory can be conserved by summarizing and filtering routes in the distribution router. The remote router need not receive routes that have been learned from other networks because the remote router must send all nonlocal traffic, regardless of destination, to the distribution router. If a true stub network is desired, the distribution router should be configured to send only a default route to the remote router. The EIGRP Stub Routing feature does not automatically enable summarization on the distribution router. In most cases, the network administrator needs to configure summarization on the distribution routers.

Without the stub feature, even after the routes that are sent from the distribution router to the remote router have been filtered or summarized, a problem might occur. If a route is lost somewhere in the corporate network, EIGRP could send a query to the distribution router, which in turn sends a query to the remote router even if routes are being summarized. If there is a problem communicating over the WAN link between the distribution router and the remote router, an EIGRP stuck in active (SIA) condition could occur and cause instability elsewhere in the network. The EIGRP Stub Routing feature allows a network administrator to prevent queries from being sent to the remote router.

## Route Policy Options for an EIGRP Process

Route policies comprise series of statements and expressions that are bracketed with the **route-policy** and **end-policy** keywords. Rather than a collection of individual commands (one for each line), the statements within a route policy have context relative to each other. Thus, instead of each line being an individual command, each policy or set is an independent configuration object that can be used, entered, and manipulated as a unit.

Each line of a policy configuration is a logical subunit. At least one new line must follow the **then**, **else**, and **end-policy** keywords. A new line must also follow the closing parenthesis of a parameter list and the name string in a reference to an AS path set, community set, extended community set, or prefix set (in the EIGRP context). At least one new line must precede the definition of a route policy or prefix set. A new line must appear at the end of a logical unit of policy expression and may not appear anywhere else.



This is the command to set the EIGRP metric in a route policy:

```
RP/0/RP0/CPU0:router(config-rpl)# set eigrp-metric bandwidth delay reliability loading mtu
```

This is the command to provide EIGRP offset list functionality in a route policy:

```
RP/0/RP0/CPU0:router(config-rpl)# add eigrp-metric bandwidth delay reliability loading mtu
```

A route policy can be used in EIGRP only if all the statements are applicable to the particular EIGRP attach point. The following commands accept a route policy:

- **default-information allowed**—Match statements are allowed for destination. No set statements are allowed.
- **route-policy**—Match statements are allowed for destination, next hop, and tag. Set statements are allowed for eigrp-metric and tag.
- **redistribute**—Match statements are allowed for destination, next hop, source-protocol, tag and route-type. Set statements are allowed for eigrp-metric and tag.

The range for setting a tag is 0 to 255 for internal routes and 0 to 4294967295 for external routes.

## EIGRP Layer 3 VPN PE-CE Site-of-Origin

The EIGRP MPLS and IP VPN PE-CE Site-of-Origin (SoO) feature introduces the capability to filter Multiprotocol Label Switching (MPLS) and IP Virtual Private Network (VPN) traffic on a per-site basis for EIGRP networks. SoO filtering is configured at the interface level and is used to manage MPLS and IP VPN traffic and to prevent transient routing loops from occurring in complex and mixed network topologies.

### Router Interoperation with the Site-of-Origin Extended Community

The configuration of the SoO extended community allows routers that support this feature to identify the site from which each route originated. When this feature is enabled, the EIGRP routing process on the PE or CE router checks each received route for the SoO extended community and filters based on the following conditions:

- A received route from BGP or a CE router contains a SoO value that matches the SoO value on the receiving interface:
  - If a route is received with an associated SoO value that matches the SoO value that is configured on the receiving interface, the route is filtered out because it was learned from another PE router or from a backdoor link. This behavior is designed to prevent routing loops.
- A received route from a CE router is configured with a SoO value that does not match:
  - If a route is received with an associated SoO value that does not match the SoO value that is configured on the receiving interface, the route is accepted into the EIGRP topology table so that it can be redistributed into BGP.
  - If the route is already installed in the EIGRP topology table but is associated with a different SoO value, the SoO value from the topology table is used when the route is redistributed into BGP.
- A received route from a CE router does not contain a SoO value:
  - If a route is received without a SoO value, the route is accepted into the EIGRP topology table, and the SoO value from the interface that is used to reach the next-hop CE router is appended to the route before it is redistributed into BGP.

When BGP and EIGRP peers that support the SoO extended community receive these routes, they also receive the associated SoO values and pass them to other BGP and EIGRP peers that support the SoO extended community. This filtering is designed to prevent transient routes from being relearned from the originating site, which prevents transient routing loops from occurring.

In conjunction with BGP cost community, EIGRP, BGP, and the RIB ensure that paths over the MPLS VPN core are preferred over backdoor links.

For MPLS and IP VPN and SoO configuration information, see *Implementing MPLS Layer 3 VPNs* in *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## IPv6 and IPv6 VPN Provider Edge Support over MPLS and IP

IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) uses the existing IP and Multiprotocol Label Switching (MPLS) IPv4 core infrastructure for IPv6 transport. 6PE and 6VPE enable IPv6 sites to communicate with each other over an IP and MPLS IPv4 core network using MPLS label switched paths (LSPs).

**Note**

---

This feature is supported on Cisco XR 12000 Series Routers.

---

EIGRP is an Interior Gateway Protocol (IGP) that supports the 6PE or 6VPE provider edge-to-customer edge protocol by supporting the configuration of IPv6 address families in EIGRP VRF and exchanging IPv6 routing updates in the L3VPN environment

For detailed information on configuring 6PE and 6VPE over MPLS and IP, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## How to Implement EIGRP on Cisco IOS XR Software

This section contains instructions for the following tasks:

- [Enabling EIGRP Routing, page RC-149](#) (required)
- [Configuring Route Summarization for an EIGRP Process, page RC-151](#) (optional)
- [Redistributing Routes for EIGRP, page RC-153](#) (optional)
- [Creating a Route Policy and Attaching It to an EIGRP Process, page RC-155](#) (optional)
- [Configuring Stub Routing for an EIGRP Process, page RC-158](#) (optional)
- [Configuring EIGRP as a PE-CE Protocol, page RC-159](#) (optional)
- [Redistributing BGP Routes into EIGRP, page RC-161](#) (optional)
- [Monitoring EIGRP Routing, page RC-163](#) (optional)
- [Monitoring EIGRP Routing, page RC-163](#) (optional)

**Note**

---

To save configuration changes, you must commit changes when the system prompts you.

---

## Enabling EIGRP Routing

This task enables EIGRP routing and establishes an EIGRP routing process.

### Prerequisites

Although you can configure EIGRP before you configure an IP address, no EIGRP routing occurs until at least one IP address is configured.

### SUMMARY STEPS

1. **configure**
2. **router eigrp** *as-number*
3. **address-family** {*ipv4* | *ipv6*}
4. **router-id** *id*
5. **default-metric** *bandwidth delay reliability loading mtu*
6. **distance** *internal-distance external-distance*
7. **interface** *type instance*
8. **holdtime** *seconds*
9. **bandwidth-percent** *percent*
10. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> }	Enters an address family configuration mode.
	<b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4	

	Command or Action	Purpose
Step 4	<p><b>router-id</b> <i>id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# router-id 172.20.1.1</p>	<p>(Optional) Configures a router-id for an EIGRP process.</p> <p><b>Note</b> It is good practice to use the <b>router-id</b> command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that EIGRP can function regardless of the interface address configuration.</p>
Step 5	<p><b>default-metric</b> <i>bandwidth delay reliability loading mtu</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# default-metric 1000 100 250 100 1500</p>	<p>(Optional) Sets metrics for an EIGRP process.</p>
Step 6	<p><b>distance</b> <i>internal-distance external-distance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# distance 80 130</p>	<p>(Optional) Allows the use of two administrative distances—internal and external—that could be a better route to a node.</p>
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# interface pos 0/1/0/0</p>	<p>Defines the interfaces on which the EIGRP routing protocol runs.</p>
Step 8	<p><b>holdtime</b> <i>seconds</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af-if)# holdtime 30</p>	<p>(Optional) Configures the hold time for an interface.</p> <p><b>Note</b> To ensure nonstop forwarding during RP failovers, as the number of neighbors increase, a higher holdtime than the default value is recommended. With 256 neighbors across all VRFs, we recommend 60 seconds.</p>

	Command or Action	Purpose
Step 9	<b>bandwidth-percent</b> <i>percent</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af-if)# bandwidth-percent 75	(Optional) Configures the percentage of bandwidth that may be used by EIGRP on an interface.
Step 10	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af-if)# end OR RP/0/RP0/CPU0:router(config-eigrp-af-if)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Route Summarization for an EIGRP Process

This task configures route summarization for an EIGRP process.

You can configure a summary aggregate address for a specified interface. If any more specific routes are in the routing table, EIGRP advertises the summary address from the interface with a metric equal to the minimum of all more specific routes.

### Restrictions

You should not use the **summary-address** summarization command to generate the default route (0.0.0.0) from an interface. This command creates an EIGRP summary default route to the null 0 interface with an administrative distance of 5. The low administrative distance of this default route can cause this route to displace default routes learned from other neighbors from the routing table. If the default route learned from the neighbors is displaced by the summary default route or the summary route is the only default route present, all traffic destined for the default route does not leave the router; instead, this traffic is sent to the null 0 interface, where it is dropped.

The recommended way to send only the default route from a given interface is to use a **route-policy** command.

### SUMMARY STEPS

1. **configure**

2. **router eigrp** *as-number*
3. **address-family** {*ipv4* | *ipv6*}
4. **route-policy** *name out*
5. **interface** *type instance*
6. **summary-address** *ip-address* [*length* | *mask*] [*admin-distance*]
7. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> }	Enters an address family configuration mode.
	<b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4	
Step 4	<b>route-policy</b> <i>name out</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# route-policy FILTER_DEFAULT out	Applies a routing policy to updates advertised to or received from an EIGRP neighbor.
Step 5	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# interface pos 0/1/0/0	Defines the interfaces on which the EIGRP routing protocol runs.

	Command or Action	Purpose
Step 6	<p><b>summary-address</b> <i>ip-address</i> [/length   mask] [<i>admin-distance</i>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af-if)# summary-address 192.168.0.0/16 95</p>	Configures a summary aggregate address for the specified EIGRP interface.
Step 7	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af-if)# end Or RP/0/RP0/CPU0:router(config-eigrp-af-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing Routes for EIGRP

This task explains how to redistribute routes, apply limits on the number of routes, and set timers for nonstop forwarding.

### SUMMARY STEPS

1. **configure**
2. **router eigrp** *as-number*
3. **address-family**{*ipv4* | *ipv6*}
4. **redistribute** { {**bgp** | **connected** | **isis** | **ospf** | **ospfv3** | **rip** | **static** } [*as-number*] } [**route-policy** *name*]
5. **redistribute maximum-prefix** *maximum* [*threshold*] [[**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] | [**warning-only**]]
6. **timers nsf route-hold** *seconds*
7. **maximum paths** *maximum*
8. **maximum-prefix** *maximum* [*threshold*] [[**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] | [**warning-only**]]

```

9. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4	Enters an address family configuration mode.
Step 4	<b>redistribute</b> {( <i>bgp</i>   <i>connected</i>   <i>isis</i>   <i>ospf</i>   <i>ospfv3</i>   <i>rip</i>   <i>static</i> ) [ <i>as-number</i> ]} [ <i>route-policy name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# redistribute bgp 100	Injects routes from one routing domain into EIGRP.
Step 5	<b>redistribute maximum-prefix</b> <i>maximum</i> [ <i>threshold</i> ] [[ <i>dampened</i> ] [ <i>reset-time minutes</i> ]] [ <i>restart minutes</i> ] [ <i>restart-count number</i> ]   [ <i>warning-only</i> ]]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# redistribute maximum-prefix 5000 95 warning-only	Limits the number of prefixes redistributed into an EIGRP process.
Step 6	<b>timers nsf route-hold</b> <i>seconds</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# timers nsf route-hold 120	Sets the timer that determines how long an NSF-aware EIGRP router holds routes for an inactive peer.
Step 7	<b>maximum paths</b> <i>maximum</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# maximum paths 10	Controls the maximum number of parallel routes that the EIGRP can support.



	Command or Action	Purpose
Step 8	<pre>maximum-prefix maximum [threshold] [[dampened] [reset-time minutes] [restart minutes] [restart-count number]   [warning-only]]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# maximum-prefix 50000</p>	Limits the number of prefixes that are accepted under an address family by EIGRP.
Step 9	<pre>end or commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# end OR RP/0/RP0/CPU0:router(config-eigrp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Creating a Route Policy and Attaching It to an EIGRP Process

This task defines a route policy and shows how to attach it to an EIGRP process.

A route policy definition consists of the **route-policy** command and *name* argument followed by a sequence of optional policy statements, and then closed with the **end-policy** command.

A route policy is not useful until it is applied to routes of a routing protocol.

### SUMMARY STEPS

1. **configure**
2. **route-policy** *name*
3. **set eigrp-metric** *bandwidth delay reliability load mtu*
4. **end-policy**
5. **end**  
or  
**commit**
6. **configure**
7. **router eigrp** *as-number*

8. **address-family** {ipv4 | ipv6}
9. **route-policy** *route-policy-name* {in | out}
10. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy IN-IPv4	Defines a route policy and enters route-policy configuration mode.
Step 3	<b>set eigrp-metric</b> <i>bandwidth delay reliability load mtu</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# set eigrp metric 42 100 200 100 1200	(Optional) Sets the EIGRP metric attribute.
Step 4	<b>end-policy</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end-policy	Ends the definition of a route policy and exits route-policy configuration mode.
Step 5	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end or RP/0/RP0/CPU0:router(config-rpl)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

	Command or Action	Purpose
Step 6	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 7	<p><b>router eigrp</b> <i>as-number</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100</p>	Configures an EIGRP routing process.
Step 8	<p><b>address-family</b> {<b>ipv4</b>   <b>ipv6</b>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4</p>	Enters an address family configuration mode.
Step 9	<p><b>route-policy</b> <i>route-policy-name</i> {<b>in</b>   <b>out</b>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# route-policy IN-IPv4 in</p>	Applies a routing policy to updates advertised to or received from an EIGRP neighbor.
Step 10	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# end or RP/0/RP0/CPU0:router(config-eigrp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Stub Routing for an EIGRP Process

This task configures the distribution and remote routers to use an EIGRP process for stub routing.

### Restrictions

EIGRP stub routing should be used only on remote routers. A stub router is defined as a router connected to the network core or distribution layer through which core transit traffic should not flow. A stub router should not have any EIGRP neighbors other than distribution routers. Ignoring this restriction causes undesirable behavior.

### SUMMARY STEPS

1. **configure**
2. **router eigrp** *as-number*
3. **address-family** {*ipv4* | *ipv6*}
4. **stub** [*receive-only* | {[*connected*] [*redistributed*] [*static*] [*summary*]}]
5. **end**  
or  
**commit**
6. **show eigrp** [*ipv4* | *ipv6*] [*vrf* {*vrf-name* | *all*}] **neighbors** [*as-number*] [*detail*] [*interface-type* *interface-instance* | *static*]

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> }	Enters an address family configuration mode.
Step 4	<b>stub</b> [ <i>receive-only</i>   {[ <i>connected</i> ] [ <i>redistributed</i> ] [ <i>static</i> ] [ <i>summary</i> ]}]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# stub receive-only	Configures a router as a stub for EIGRP.

	Command or Action	Purpose
Step 5	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-eigrp-af)# end  OR  RP/0/RP0/CPU0:router(config-eigrp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 6	<pre>show eigrp [ipv4   ipv6] [vrf {vrf-name   all}] neighbors [as-number] [detail] [interface-type interface-instance   static]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show eigrp neighbors  detail</p>	<p>Verifies that a remote router has been configured as a stub router with EIGRP.</p> <p>The last line of the output shows the stub status of the remote or spoke router.</p>

## Configuring EIGRP as a PE-CE Protocol

Perform this task to configure EIGRP on the provider edge (PE) and establish provider edge-to-customer edge (PE-CE) communication using EIGRP.

### SUMMARY STEPS

- configure**
- router eigrp** *as-number*
- vrf** *vrf-name*
- address-family** {*ipv4* | *ipv6*}
- router-id** *router-id*
- autonomous-system** *as-number*
- redistribute** {{**bgp** | **connected** | **isis** | **ospf** | **ospfv3** | **rip** | **static**} [*as-number* | *instance-name*]} [*route-policy name*]
- interface** *type instance*
- site-of-origin** {*as-number:number* | *ip-address:number*}

```

10. end
    or
    commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# vrf vrf_A	Configures a VPN routing and forwarding (VRF) instance.
Step 4	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf)# address-family ipv4	Enters a VRF address family configuration mode.
Step 5	<b>router-id</b> <i>router-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# router-id 33	Configures a router ID for the EIGRP process.
Step 6	<b>autonomous-system</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# autonomous-system 2	Configures an EIGRP routing process to run within the VRF instance.
Step 7	<b>redistribute</b> {{ <b>bgp</b>   <b>connected</b>   <b>isis</b>   <b>ospf</b>   <b>ospfv3</b>   <b>rip</b>   <b>static</b> } [ <i>as-number</i> ]} [ <i>route-policy name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# redistribute bgp 100	Injects routes from one routing domain into EIGRP.
Step 8	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# interface gigabitEthernet 0/1/5/0	Configures the interface on which EIGRP the routing protocol runs.

	Command or Action	Purpose
Step 9	<pre>site-of-origin {as-number:number   ip-address:number}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af-if)# site-of-origin 3:4 </p>	Configures the site-of-origin (SoO) filtering on the EIGRP interface.
Step 10	<pre>end or commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af-if)# end OR RP/0/RP0/CPU0:router(config-eigrp-vrf-af-if)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing BGP Routes into EIGRP

Perform this task to redistribute BGP routes into EIGRP.

Typically, EIGRP routes are redistributed into BGP with extended community information appended to the route. BGP carries the route over the VPN backbone with the EIGRP-specific information encoded in the BGP extended community attributes. After the peering customer site receives the route, EIGRP redistributes the BGP route then extracts the BGP extended community information and reconstructs the route as it appeared in the original customer site.

When redistributing BGP routes into EIGRP, the receiving provider edge (PE) EIGRP router looks for BGP extended community information. If the information is received, it is used to recreate the original EIGRP route. If the information is missing, EIGRP uses the configured default metric value.

If the metric values are not derived from the BGP extended community and a default metric is not configured, the route is not advertised to the customer edge (CE) router by the PE EIGRP. When BGP is redistributed into BGP, metrics may not be added to the BGP prefix as extended communities; for example, if EIGRP is not running on the other router. In this case, EIGRP is redistributed into BGP with a “no-metrics” option.

### SUMMARY STEPS

1. **configure**
2. **router eigrp** *as-number*

3. **vrf** *vrf-name*
4. **address-family** {**ipv4** | **ipv6**}
5. **redistribute** {{**bgp** | **connected** | **isis** | **ospf** | **ospfv3** | **rip** | **static**} [*as-number* | *instance-name*]} [*route-policy name*]
6. **route-policy** *route-policy-name* {**in** | **out**}
7. **default-metric** *bandwidth delay reliability loading mtu*
8. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# router eigrp 100	Configures a VRF instance.
Step 4	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> }	Enters a VRF address family configuration mode.
Step 5	<b>redistribute</b> {{ <b>bgp</b>   <b>connected</b>   <b>isis</b>   <b>ospf</b>   <b>ospfv3</b>   <b>rip</b>   <b>static</b> } [ <i>as-number</i> ]} [ <i>route-policy name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# redistribute bgp 100	Injects routes from one routing domain into EIGRP.
Step 6	<b>route-policy</b> <i>route-policy-name</i> { <b>in</b>   <b>out</b> }	Applies a routing policy to updates advertised to or received from an EIGRP neighbor.
	<b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# route-policy policy_A in	



	Command or Action	Purpose
Step 7	<pre>default-metric bandwidth delay reliability loading mtu</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# default-metric 1000 100 250 100 1500</pre>	Configures metrics for EIGRP.
Step 8	<pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# end Or RP/0/RP0/CPU0:router(config-eigrp-vrf-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Monitoring EIGRP Routing

The commands in this section are used to log neighbor adjacency changes, monitor the stability of the routing system, and help detect problems.

### SUMMARY STEPS

1. **configure**
2. **router eigrp** *as-number*
3. **address-family** [*ipv4* | *ipv6*]
4. **log-neighbor-changes**
5. **log-neighbor-warnings**
6. **end**  
or  
**commit**
7. **clear eigrp** [*as-number*] [*vrf {vrf | all}*] [*ipv4* | *ipv6*] **neighbors** [*ip-address* | *interface-type interface-instance*]
8. **clear eigrp** [*as-number*] [*vrf {vrf | all}*] [*ipv4* | *ipv6*] **topology** [*prefix mask*] [*prefixlength*]
9. **show eigrp** [*as-number*] [*vrf {vrf | all}*] [*ipv4* | *ipv6*] **accounting**

10. **show eigrp** [*as-number*] [**vrf** {*vrf* | **all**}] [**ipv4** | **ipv6**] **interfaces** [*type instance*] [**detail**]
11. **show eigrp** [*as-number*] [**vrf** {*vrf* | **all**}] [**ipv4** | **ipv6**]**neighbors** [**detail**] [*interface-type interface-instance* | **static**]
12. **show protocols eigrp** [**vrf** *vrf-name*]
13. **show eigrp** [*as-number*] [**vrf** {*vrf* | **all**}] [**ipv4** | **ipv6**] **topology** [*ip-address mask*] [**active** | **all-links** | **detail-links** | **pending** | **summary** | **zero-successors**]
14. **show eigrp** [*as-number*] [**vrf** {*vrf* | **all**}] [**ipv4** | **ipv6**] **traffic**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router eigrp</b> <i>as-number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router eigrp 100	Configures an EIGRP routing process.
Step 3	<b>address-family</b> [ <b>ipv4</b>   <b>ipv6</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp)# <b>address-family ipv4</b>	Enters an address family configuration mode.
Step 4	<b>log-neighbor-changes</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# log-neighbor-changes	Enables the logging of changes in EIGRP neighbor adjacencies.
Step 5	<b>log-neighbor-warnings</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-eigrp-af)# log-neighbor-warnings	Enables the logging of EIGRP neighbor warning messages.

	Command or Action	Purpose
Step 6	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-eigrp-af)# end  OR  RP/0/RP0/CPU0:router(config-eigrp-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)?  [cancel]:</li> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 7	<pre>clear eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] neighbors [ip-address   interface-type interface-instance]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# clear eigrp 20 neighbors pos 0/1/0/0</p>	<p>Deletes EIGRP and VPN neighbor entries from the appropriate table.</p>
Step 8	<pre>clear eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] topology [prefix mask] [prefix/length]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# clear eigrp topology</p>	<p>Deletes EIGRP and VRF topology entries from the appropriate tab.</p>
Step 9	<pre>show eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] accounting</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show eigrp all accounting</p>	<p>Displays prefix accounting information for EIGRP processes.</p>
Step 10	<pre>show eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] interfaces [type instance] [detail]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show eigrp interfaces detail</p>	<p>Displays information about interfaces configured for EIGRP.</p>

	Command or Action	Purpose
Step 11	<pre>show eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] neighbors [detail] [interface-type interface-instance   static]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show eigrp neighbors 20 detail static</p>	Displays the neighbors discovered by EIGRP.
Step 12	<pre>show protocols eigrp [vrf vrf-name]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show protocols eigrp</p>	Displays information about the EIGRP process configuration.
Step 13	<pre>show eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] topology [ip-address mask] [active   all-links   detail-links   pending   summary   zero-successors]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show eigrp topology 10.0.0.1 253.254.255.255 summary</p>	Displays entries in the EIGRP topology table.
Step 14	<pre>show eigrp [as-number] [vrf {vrf   all}] [ipv4   ipv6] traffic</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show eigrp traffic</p>	Displays the number of EIGRP packets sent and received.

## Configuration Examples for Implementing EIGRP on Cisco IOS XR Software

This section provides the following configuration examples:

- [Configuring a Basic EIGRP Configuration: Example, page RC-166](#)
- [Configuring an EIGRP Stub Operation: Example, page RC-167](#)
- [Configuring an EIGRP PE-CE Configuration with Prefix-Limits: Example, page RC-167](#)

### Configuring a Basic EIGRP Configuration: Example

The following example shows how to configure EIGRP with a policy that filters incoming routes. This is a typical configuration for a router that has just one neighbor, but advertises other connected subnets.

```
router eigrp 144
 address-family ipv4
  metric maximum-hops 20
  router-id 102.10.9.4
  route-policy GLOBAL_FILTER_POLICY in
  log-neighbor-changes
  log-neighbor-warnings
  interface Loopback0
  !
```

```
interface POS0/2/0/0
  passive-interface
!
interface GigabitEthernet0/6/0/0
  hello-interval 8
  hold-time 30
  summary-address 10.0.0.0 255.255.0.0
!
!
```

## Configuring an EIGRP Stub Operation: Example

The following example shows how to configure an EIGRP stub. Stub operation allows only connected, static, and summary routes to be advertised to neighbors.

```
router eigrp 200
  address-family ipv4
    stub connected static summary
    router-id 172.16.82.22
    log-neighbor-changes
    log-neighbor-warnings
    redistribute connected route-policy CONN_POLICY
  interface GigabitEthernet0/6/0/0
    passive-interface
    neighbor 10.0.0.31
  !
  interface GigabitEthernet0/6/0/1
    passive-interface
    neighbor 10.0.1.21
  !
!
```

## Configuring an EIGRP PE-CE Configuration with Prefix-Limits: Example

The following example shows how to configure EIGRP to operate as a PE-CE protocol on a PE router. The configuration is under VRF CUSTOMER\_1. A maximum prefix is typically configured to ensure that one set of customer routes do not overwhelm the EIGRP process.

```
router eigrp 500
  vrf CUSTOMER_1
    address-family ipv4
      timers nsf route-hold 300
      router-id 172.16.6.11
      maximum-prefix 450 70
      default-metric 200000 10000 195 10 1500
      log-neighbor-changes
      log-neighbor-warnings
      redistribute maximum-prefix 350 70
      redistribute bgp 1.65500 route-policy SITE_1_POLICY
    interface POS0/4/0/5
      neighbor 10.22.1.1
    !
  !
!
```

# Additional References

The following sections provide references related to implementing EIGRP on Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
EIGRP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference</i>
MPLS VPN support for EIGRP feature information	<i>Implementing MPLS Layer 3 VPNs</i> module and <i>Implementing MPLS Layer 2 VPNs</i> module in the <i>Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>
Site of Origin (SoO) support for EIGRP feature information	<i>Implementing MPLS Traffic Engineering on Cisco IOS XR Software</i> module in the <i>Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>

## Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing standards has not been modified by this feature.	—

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>







# Implementing OSPF on Cisco IOS XR Software

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) developed by the OSPF working group of the Internet Engineering Task Force (IETF). Designed expressly for IP networks, OSPF supports IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets.

OSPF Version 3 (OSPFv3) expands on OSPF Version 2, providing support for IPv6 routing prefixes.

This module describes the concepts and tasks you need to implement both versions of OSPF on your Cisco IOS XR router. The term “OSPF” implies both versions of the routing protocol, unless otherwise noted.



## Note

For more information about OSPF on the Cisco IOS XR software and complete descriptions of the OSPF commands listed in this module, see the “[Related Documents](#)” section of this module. To locate documentation for other commands that might appear during execution of a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing OSPF on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router.
Release 3.2.2	Support was added for OSPFv3 Graceful Restart.
Release 3.3.0	Support was added for the following features: <ul style="list-style-type: none"><li>• Multicast-Intact for OSPFv2</li><li>• Interface Association to a VRF</li><li>• OSPF Provider Edge to Customer Edge (PE-CE) Protocol</li><li>• Multiple OSPF Instances (OSPF Process and a VRF)</li><li>• RPL-based Type 3 Filtering</li><li>• LSA Pacing</li></ul>
Release 3.4.0	Support was added for the following features: <ul style="list-style-type: none"><li>• OSPF Forwarding Adjacency</li><li>• OSPF SNMP Trap MIB</li></ul>
Release 3.4.1	Support was added for the multi-area adjacency feature.

Release 3.5.0	Support was added for the following features: <ul style="list-style-type: none"> <li>• <a href="#">Label Distribution Protocol IGP Auto-configuration for OSPF</a></li> <li>• <a href="#">OSPF Authentication Message Digest Management</a></li> <li>• <a href="#">GTSM TTL Security Mechanism for OSPF</a></li> <li>• <a href="#">Path Computation Element for OSPFv2</a></li> <li>• <a href="#">OSPF Warm Standby</a></li> </ul>
Release 3.6.0	Support was added for the following features: <ul style="list-style-type: none"> <li>• <a href="#">OSPFv2 Shamlink Support for MPLS VPN</a></li> <li>• <a href="#">OSPFv2 nonstop routing (NSR)</a></li> </ul>

## Contents

- [Prerequisites for Implementing OSPF on Cisco IOS XR Software, page RC-172](#)
- [Information About Implementing OSPF on Cisco IOS XR Software, page RC-173](#)
- [How to Implement OSPF on Cisco IOS XR Software, page RC-194](#)
- [Configuration Examples for Implementing OSPF on Cisco IOS XR Software, page RC-255](#)
- [Where to Go Next, page RC-262](#)
- [Additional References, page RC-262](#)

## Prerequisites for Implementing OSPF on Cisco IOS XR Software

The following are prerequisites for implementing OSPF on Cisco IOS XR Software:

- You must be in a user group associated with a task group that includes the proper task IDs for OSPF commands. Task IDs for commands are listed in the Cisco IOS XR Task ID Reference Guide. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of the *Cisco IOS XR System Security Configuration Guide*.
- Configuration tasks for OSPFv3 assume that you are familiar with IPv6 addressing and basic configuration. See the *Implementing Network Stack IPv4 and IPv6 on Cisco IOS XR Software* module of the *Cisco IOS XR IP Addresses and Services Configuration Guide* for information on IPv6 routing and addressing.
- Before you enable OSPFv3 on an interface, you must perform the following tasks:
  - Complete the OSPF network strategy and planning for your IPv6 network. For example, you must decide whether multiple areas are required.
  - Enable IPv6 on the interface.
- Configuring authentication (IP Security) is an optional task. If you choose to configure authentication, you must first decide whether to configure plain text or Message Digest 5 (MD5) authentication, and whether the authentication applies to an entire area or specific interfaces.

# Information About Implementing OSPF on Cisco IOS XR Software

To implement OSPF you need to understand the following concepts:

- [OSPF Functional Overview](#), page RC-173
- [Key Features Supported in the Cisco IOS XR OSPF Implementation](#), page RC-175
- [Comparison of Cisco IOS XR OSPFv3 and OSPFv2](#), page RC-175
- [OSPF Hierarchical CLI and CLI Inheritance](#), page RC-176
- [OSPF Routing Components](#), page RC-176
- [OSPF Process and Router ID](#), page RC-179
- [Supported OSPF Network Types](#), page RC-179
- [Route Authentication Methods for OSPF](#), page RC-180
- [Neighbors and Adjacency for OSPF](#), page RC-181
- [Designated Router \(DR\) for OSPF](#), page RC-181
- [Default Route for OSPF](#), page RC-181
- [Link-State Advertisement Types for OSPF Version 2](#), page RC-181
- [Link-State Advertisement Types for OSPFv3](#), page RC-182
- [Virtual Link and Transit Area for OSPF](#), page RC-183
- [OSPFv2 Shamlink Support for MPLS VPN](#), page RC-184
- [Route Redistribution for OSPF](#), page RC-186
- [OSPF Shortest Path First Throttling](#), page RC-186
- [Nonstop Forwarding for OSPF Version 2](#), page RC-187
- [Graceful Restart for OSPFv3](#), page RC-188
- [Warm Standby and Nonstop Routing for OSPF Version 2](#), page RC-190
- [Multicast-Intact Support for OSPF](#), page RC-190
- [Load Balancing in OSPF Version 2 and OSPFv3](#), page RC-191
- [Multi-Area Adjacency for OSPF Version 2](#), page RC-191
- [Label Distribution Protocol IGP Auto-configuration for OSPF](#), page RC-192
- [OSPF Authentication Message Digest Management](#), page RC-193
- [GTSM TTL Security Mechanism for OSPF](#), page RC-193
- [Path Computation Element for OSPFv2](#), page RC-193

## OSPF Functional Overview

OSPF is a routing protocol for IP. It is a link-state protocol, as opposed to a distance-vector protocol. A link-state protocol makes its routing decisions based on the states of the links that connect source and destination machines. The state of the link is a description of that interface and its relationship to its

neighboring networking devices. The interface information includes the IP address of the interface, network mask, type of network to which it is connected, routers connected to that network, and so on. This information is propagated in various types of link-state advertisements (LSAs).

A router stores the collection of received LSA data in a link-state database. This database includes LSA data for the links of the router. The contents of the database, when subjected to the Dijkstra algorithm, extract data to create an OSPF routing table. The difference between the database and the routing table is that the database contains a complete collection of raw data; the routing table contains a list of shortest paths to known destinations through specific router interface ports.

OSPF is the IGP of choice because it scales to large networks. It uses areas to partition the network into more manageable sizes and to introduce hierarchy in the network. A router is attached to one or more areas in a network. All of the networking devices in an area maintain the same complete database information about the link states in their area only. They do not know about all link states in the network. The agreement of the database information among the routers in the area is called convergence.

At the intradomain level, OSPF can import routes learned using Intermediate System-to-Intermediate System (IS-IS). OSPF routes can also be exported into IS-IS. At the interdomain level, OSPF can import routes learned using Border Gateway Protocol (BGP). OSPF routes can be exported into BGP.

Unlike Routing Information Protocol (RIP), OSPF does not provide periodic routing updates. On becoming neighbors, OSPF routers establish an adjacency by exchanging and synchronizing their databases. After that, only changed routing information is propagated. Every router in an area advertises the costs and states of its links, sending this information in an LSA. This state information is sent to all OSPF neighbors one hop away. All the OSPF neighbors, in turn, send the state information unchanged. This flooding process continues until all devices in the area have the same link-state database.

To determine the best route to a destination, the software sums all of the costs of the links in a route to a destination. After each router has received routing information from the other networking devices, it runs the shortest path first (SPF) algorithm to calculate the best path to each destination network in the database.

The networking devices running OSPF detect topological changes in the network, flood link-state updates to neighbors, and quickly converge on a new view of the topology. Each OSPF router in the network soon has the same topological view again. OSPF allows multiple equal-cost paths to the same destination. Since all link-state information is flooded and used in the SPF calculation, multiple equal cost paths can be computed and used for routing.

On broadcast and nonbroadcast multiaccess (NBMA) networks, the designated router (DR) or backup DR performs the LSA flooding. On point-to-point networks, flooding simply exits an interface directly to a neighbor.

OSPF runs directly on top of IP; it does not use TCP or User Datagram Protocol (UDP). OSPF performs its own error correction by means of checksums in its packet header and LSAs.

In OSPFv3, the fundamental concepts are the same as OSPF Version 2, except that support is added for the increased address size of IPv6. New LSA types are created to carry IPv6 addresses and prefixes, and the protocol runs on an individual link basis rather than on an individual IP-subnet basis.

OSPF typically requires coordination among many internal routers: Area Border Routers (ABRs), which are routers attached to multiple areas, and Autonomous System Border Routers (ASBRs) that export reroutes from other sources (for example, IS-IS, BGP, or static routes) into the OSPF topology. At a

minimum, OSPF-based routers or access servers can be configured with all default parameter values, no authentication, and interfaces assigned to areas. If you intend to customize your environment, you must ensure coordinated configurations of all routers.

## Key Features Supported in the Cisco IOS XR OSPF Implementation

The Cisco IOS XR implementation of OSPF conforms to the OSPF Version 2 and OSPF Version 3 specifications detailed in the Internet RFC 2328 and RFC 2740, respectively.

The following key features are supported in the Cisco IOS XR implementation:

- Hierarchy—CLI hierarchy is supported.
- Inheritance—CLI inheritance is supported.
- Stub areas—Definition of stub areas is supported.
- NSF—Nonstop forwarding is supported.
- SPF throttling—Shortest path first throttling feature is supported.
- LSA throttling—LSA throttling feature is supported.
- Fast convergence—SPF and LSA throttle timers are set, configuring fast convergence. The OSPF LSA throttling feature provides a dynamic mechanism to slow down LSA updates in OSPF during network instability. LSA throttling also allows faster OSPF convergence by providing LSA rate limiting in milliseconds.
- Route redistribution—Routes learned using any IP routing protocol can be redistributed into any other IP routing protocol.
- Authentication—Plain text and MD5 authentication among neighboring routers within an area is supported.
- Routing interface parameters—Configurable parameters supported include interface output cost, retransmission interval, interface transmit delay, router priority, router “dead” and hello intervals, and authentication key.
- Virtual links—Virtual links are supported.
- Not-so-stubby area (NSSA)—RFC 1587 is supported.
- OSPF over demand circuit—RFC 1793 is supported.

## Comparison of Cisco IOS XR OSPFv3 and OSPFv2

Much of the OSPFv3 protocol is the same as in OSPFv2. OSPFv3 is described in RFC 2740.

The key differences between the Cisco IOS XR OSPFv3 and OSPFv2 protocols are as follows:

- OSPFv3 expands on OSPFv2 to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses.
- When using an NBMA interface in OSPFv3, users must manually configure the router with the list of neighbors. Neighboring routers are identified by the link local address of the attached interface of the neighbor.
- Unlike in OSPFv2, multiple OSPFv3 processes can be run on a link.
- LSAs in OSPFv3 are expressed as “prefix and prefix length” instead of “address and mask.”
- The router ID is a 32-bit number with no relationship to an IPv6 address.

## OSPF Hierarchical CLI and CLI Inheritance

Cisco IOS XR software introduces new OSPF configuration fundamentals consisting of hierarchical CLI and CLI inheritance.

Hierarchical CLI is the grouping of related network component information at defined hierarchical levels such as at the router, area, and interface levels. Hierarchical CLI allows for easier configuration, maintenance, and troubleshooting of OSPF configurations. When configuration commands are displayed together in their hierarchical context, visual inspections are simplified. Hierarchical CLI is intrinsic for CLI inheritance to be supported.

With CLI inheritance support, you need not explicitly configure a parameter for an area or interface. In Cisco IOS XR, the parameters of interfaces in the same area can be exclusively configured with a single command, or parameter values can be inherited from a higher hierarchical level—such as from the area configuration level or the router ospf configuration levels.

For example, the hello interval value for an interface is determined by this precedence “IF” statement:

If the **hello interval** command is configured at the interface configuration level, then use the interface configured value, else

If the **hello interval** command is configured at the area configuration level, then use the area configured value, else

If the **hello interval** command is configured at the router ospf configuration level, then use the router ospf configured value, else

Use the default value of the command.

**Tip**

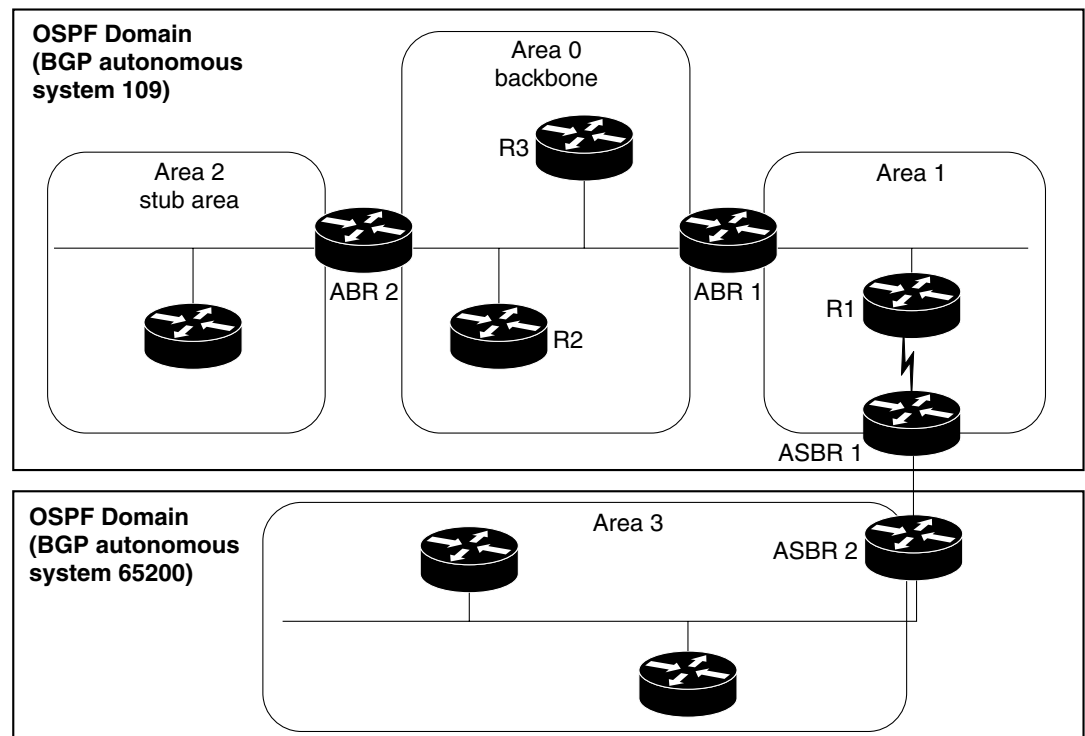
Understanding hierarchical CLI and CLI inheritance saves you considerable configuration time. See the [“Configuring Authentication at Different Hierarchical Levels for OSPF Version 2”](#) section on page RC-204 to understand how to implement these fundamentals. In addition, Cisco IOS XR examples are provided in the [“Configuration Examples for Implementing OSPF on Cisco IOS XR Software”](#) section on page RC-255.

## OSPF Routing Components

Before implementing OSPF, you must know what the routing components are and what purpose they serve. They consist of the autonomous system, area types, interior routers, ABRs, and ASBRs.

[Figure 1](#) illustrates the routing components in an OSPF network topology.

**Figure 1** OSPF Routing Components



88721

## Autonomous Systems

The autonomous system is a collection of networks, under the same administrative control, that share routing information with each other. An autonomous system is also referred to as a routing domain. [Figure 1](#) shows two autonomous systems: 109 and 65200. An autonomous system can consist of one or more OSPF areas.

## Areas

Areas allow the subdivision of an autonomous system into smaller, more manageable networks or sets of adjacent networks. As shown in [Figure 1](#), autonomous system 109 consists of three areas: Area 0, Area 1, and Area 2.

OSPF hides the topology of an area from the rest of the autonomous system. The network topology for an area is visible only to routers inside that area. When OSPF routing is within an area, it is called *intra-area routing*. This routing limits the amount of link-state information flood into the network, reducing routing traffic. It also reduces the size of the topology information in each router, conserving processing and memory requirements in each router.

Also, the routers within an area cannot see the detailed network topology outside the area. Because of this restricted view of topological information, you can control traffic flow between areas and reduce routing traffic when the entire autonomous system is a single routing domain.

## Backbone Area

A backbone area is responsible for distributing routing information between multiple areas of an autonomous system. OSPF routing occurring outside of an area is called *interarea routing*.

The backbone itself has all properties of an area. It consists of ABRs, routers, and networks only on the backbone. As shown in [Figure 1](#), Area 0 is an OSPF backbone area. Any OSPF backbone area has a reserved area ID of 0.0.0.0.

## Stub Area

A stub area is an area that does not accept or detailed network information external to the area. A stub area typically has only one router that interfaces the area to the rest of the autonomous system. The stub ABR advertises a single default route to external destinations into the stub area. Routers within a stub area use this route for destinations outside the area and the autonomous system. This relationship conserves LSA database space that would otherwise be used to store external LSAs flooded into the area. In [Figure 1](#), Area 2 is a stub area that is reached only through ABR 2. Area 0 cannot be a stub area.

## Not-so-Stubby Area

A Not-so-Stubby Area (NSSA) is similar to the stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but can import autonomous system external routes in a limited fashion within the area.

NSSA allows importing of Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA ABRs, which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

Use NSSA to simplify administration if you are a network administrator that must connect a central site using OSPF to a remote site that is using a different routing protocol.

Before NSSA, the connection between the corporate site border router and remote router could not be run as an OSPF stub area because routes for the remote site could not be redistributed into a stub area, and two routing protocols needed to be maintained. A simple protocol like RIP was usually run and handled the redistribution. With NSSA, you can extend OSPF to cover the remote connection by defining the area between the corporate router and remote router as an NSSA. Area 0 cannot be an NSSA.

## Routers

The OSPF network is composed of ABRs, ASBRs, and interior routers.

### Area Border Routers

An area border routers (ABR) is a router with multiple interfaces that connect directly to networks in two or more areas. An ABR runs a separate copy of the OSPF algorithm and maintains separate routing data for each area that is attached to, including the backbone area. ABRs also send configuration summaries for their attached areas to the backbone area, which then distributes this information to other OSPF areas in the autonomous system. In [Figure 1](#), there are two ABRs. ABR 1 interfaces Area 1 to the backbone area. ABR 2 interfaces the backbone Area 0 to Area 2, a stub area.

### Autonomous System Boundary Routers (ASBR)

An autonomous system boundary router (ASBR) provides connectivity from one autonomous system to another system. ASBRs exchange their autonomous system routing information with boundary routers in other autonomous systems. Every router inside an autonomous system knows how to reach the boundary routers for its autonomous system.



ASBRs can import external routing information from other protocols like BGP and redistribute them as AS-external (ASE) Type 5 LSAs to the OSPF network. If the Cisco IOS XR router is an ASBR, you can configure it to advertise VIP addresses for content as autonomous system external routes. In this way, ASBRs flood information about external networks to routers within the OSPF network.

ASBR routes can be advertised as a Type 1 or Type 2 ASE. The difference between Type 1 and Type 2 is how the cost is calculated. For a Type 2 ASE, only the external cost (metric) is considered when multiple paths to the same destination are compared. For a Type 1 ASE, the combination of the external cost and cost to reach the ASBR is used. Type 2 external cost is the default and is always more costly than an OSPF route and used only if no OSPF route exists.

## Interior Routers

The interior routers (such as R1 in [Figure 1](#)) attached to one area (for example, all the interfaces reside in the same area).

## OSPF Process and Router ID

An OSPF process is a logical routing entity running OSPF in a physical router. This logical routing entity should not be confused with the logical routing feature that allows a system administrator (known as the Cisco IOS XR Owner) to partition the physical box into separate routers.

A physical router can run multiple OSPF processes, although the only reason to do so would be to connect two or more OSPF domains. Each process has its own link-state database. The routes in the routing table are calculated from the link-state database. One OSPF process does not share routes with another OSPF process unless the routes are redistributed.

Each OSPF process is identified by a router ID. The router ID must be unique across the entire routing domain. OSPFv2 obtains a router ID from the following sources, in order of decreasing preference:

OSPF attempts to obtain a router ID in the following ways (in order of preference):

- The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router and need not be a routable IPv4 address.)
- The 32-bit numeric value specified by the **router-id** command in global configuration mode. (This value must be an IPv4 address assigned to an interface on this router.)
- The highest IPv4 address on a loopback interface in the system.
- The primary IPv4 address of an interface over which this OSPF process is running. The first interface address in the OSPF interface is selected.

We recommend that the router ID be set by the **router-id** command in router configuration mode. Separate OSPF processes could share the same router ID, in which case they cannot reside in the same OSPF routing domain.

## Supported OSPF Network Types

OSPF classifies different media into the following types of networks:

- NBMA networks
- Point-to-point networks (POS)
- Broadcast networks (Gigabit Ethernet)

- Point-to-multipoint

You can configure your Cisco IOS XR network as either a broadcast or an NBMA network. Using this feature, you can configure broadcast networks as NBMA networks when, for example, you have routers in your network that do not support multicast addressing.

## Route Authentication Methods for OSPF

OSPF Version 2 supports two types of authentication: plain text authentication and MD5 authentication. By default, no authentication is enabled (referred to as null authentication in RFC 2178).

OSPF Version 3 supports all types of authentication except key rollover.

### Plain Text Authentication

Plain text authentication (also known as Type 1 authentication) uses a password that travels on the physical medium and is easily visible to someone that does not have access permission and could use the password to infiltrate a network. Therefore, plain text authentication does not provide security. It might protect against a faulty implementation of OSPF or a misconfigured OSPF interface trying to send erroneous OSPF packets.

### MD5 Authentication

MD5 authentication provides a means of security. No password travels on the physical medium. Instead, the router uses MD5 to produce a message digest of the OSPF packet plus the key, which is sent on the physical medium. Using MD5 authentication prevents a router from accepting unauthorized or deliberately malicious routing updates, which could compromise your network security by diverting your traffic.

**Note**

---

MD5 authentication supports multiple keys, requiring that a key number be associated with a key.

---

Also see [“OSPF Authentication Message Digest Management” section on page RC-193](#)

### Authentication Strategies

Authentication can be specified for an entire process or area, or on an interface or a virtual link. An interface or virtual link can be configured for only one type of authentication, not both. Authentication configured for an interface or virtual link overrides authentication configured for the area or process.

If you intend for all interfaces in an area to use the same type of authentication, you can configure fewer commands if you use the **authentication** command in the area configuration submode (and specify the **message-digest** keyword if you want the entire area to use MD5 authentication). This strategy requires fewer commands than specifying authentication for each interface.

### Key Rollover

To support the changing of an MD5 key in an operational network without disrupting OSPF adjacencies (and hence the topology), a key rollover mechanism is supported. As a network administrator configures the new key into the multiple networking devices that communicate, some time exists when different devices are using both a new key and an old key. If an interface is configured with a new key, the software sends two copies of the same packet, each authenticated by the old key and new key. The software tracks

which devices start using the new key, and the software stops sending duplicate packets after it detects that all of its neighbors are using the new key. The software then discards the old key. The network administrator must then remove the old key from each the configuration file of each router.

## Neighbors and Adjacency for OSPF

Routers that share a segment (Layer 2 link between two interfaces) become neighbors on that segment. OSPF uses the hello protocol as a neighbor discovery and keep alive mechanism. The hello protocol involves receiving and periodically sending hello packets out each interface. The hello packets list all known OSPF neighbors on the interface. Routers become neighbors when they see themselves listed in the hello packet of the neighbor. After two routers are neighbors, they may proceed to exchange and synchronize their databases, which creates an adjacency. On broadcast and NBMA networks all neighboring routers have an adjacency.

## Designated Router (DR) for OSPF

On point-to-point and point-to-multipoint networks, the Cisco IOS XR software floods routing updates to immediate neighbors. No DR or backup DR (BDR) exists; all routing information is flooded to each router.

On broadcast or NBMA segments only, OSPF minimizes the amount of information being exchanged on a segment by choosing one router to be a DR and one router to be a BDR. Thus, the routers on the segment have a central point of contact for information exchange. Instead of each router exchanging routing updates with every other router on the segment, each router exchanges information with the DR and BDR. The DR and BDR relay the information to the other routers. On broadcast network segments the number of OSPF packets is further reduced by the DR and BDR sending such OSPF updates to a multicast IP address that all OSPF routers on the network segment are listening on.

The software looks at the priority of the routers on the segment to determine which routers are the DR and BDR. The router with the highest priority is elected the DR. If there is a tie, then the router with the higher router ID takes precedence. After the DR is elected, the BDR is elected the same way. A router with a router priority set to zero is ineligible to become the DR or BDR.

## Default Route for OSPF

Type 5 (ASE) LSAs are generated and flooded to all areas except stub areas. For the routers in a stub area to be able to route packets to destinations outside the stub area, a default route is injected by the ABR attached to the stub area.

The cost of the default route is 1 (default) or is determined by the value specified in the **default-cost** command.

## Link-State Advertisement Types for OSPF Version 2

Each of the following LSA types has a different purpose:

- Router LSA (Type 1)—Describes the links that the router has within a single area, and the cost of each link. These LSAs are flooded within an area only. The LSA indicates if the router can compute paths based on quality of service (QoS), whether it is an ABR or ASBR, and if it is one end of a virtual link. Type 1 LSAs are also used to advertise stub networks.

- Network LSA (Type 2)—Describes the link state and cost information for all routers attached to a multiaccess network segment. This LSA lists all the routers that have interfaces attached to the network segment. It is the job of the designated router of a network segment to generate and track the contents of this LSA.
- Summary LSA for ABRs (Type 3)—Advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks aggregated into one prefix. Only ABRs generate summary LSAs.
- Summary LSA for ASBRs (Type 4)—Advertises an ASBR and the cost to reach it. Routers that are trying to reach an external network use these advertisements to determine the best path to the next hop. ABRs generate Type 4 LSAs.
- Autonomous system external LSA (Type 5)—Redistributes routes from another autonomous system, usually from a different routing protocol into OSPF.
- Autonomous system external LSA (Type 7)—Provides for carrying external route information within an NSSA. Type 7 LSAs may be originated by and advertised throughout an NSSA. NSSAs do not receive or originate Type 5 LSAs. Type 7 LSAs are advertised only within a single NSSA. They are not flooded into the backbone area or into any other area by border routers.
- Intra-area-prefix LSAs (Type 9)—A router can originate multiple intra-area-prefix LSAs for every router or transit network, each with a unique link-state ID. The link-state ID for each intra-area-prefix LSA describes its association to either the router LSA or network LSA and contains prefixes for stub and transit networks.
- Area local scope (Type 10)—Opaque LSAs are not flooded past the borders of their associated area.
- Link-state (Type 11)—The LSA is flooded throughout the AS. The flooding scope of Type 11 LSAs are equivalent to the flooding scope of AS-external (Type 5) LSAs. Similar to Type 5 LSAs, the LSA is rejected if a Type 11 opaque LSA is received in a stub area from a neighboring router within the stub area. Type 11 opaque LSAs have these attributes:
  - LSAs are flooded throughout all transit areas.
  - LSAs are not flooded into stub areas from the backbone.
  - LSAs are not originated by routers into their connected stub areas.

## Link-State Advertisement Types for OSPFv3

Each of the following LSA types has a different purpose:

- Router LSA (Type 1)—Describes the link state and costs of a the router link to the area. These LSAs are flooded within an area only. The LSA indicates whether the router is an ABR or ASBR and if it is one end of a virtual link. Type 1 LSAs are also used to advertise stub networks. In OSPFv3, these LSAs have no address information and are network protocol independent. In OSPFv3, router interface information may be spread across multiple router LSAs. Receivers must concatenate all router LSAs originated by a given router before running the SPF calculation.
- Network LSA (Type 2)—Describes the link state and cost information for all routers attached to a multiaccess network segment. This LSA lists all OSPF routers that have interfaces attached to the network segment. Only the elected designated router for the network segment can generate and track the network LSA for the segment. In OSPFv3, network LSAs have no address information and are network-protocol-independent.

- **Interarea-prefix LSA for ABRs (Type 3)**—Advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or set of networks aggregated into one prefix. Only ABRs generate Type 3 LSAs. In OSPFv3, addresses for these LSAs are expressed as “prefix and prefix length” instead of “address and mask.” The default route is expressed as a prefix with length 0.
- **Interarea-router LSA for ASBRs (Type 4)**—Advertises an ASBR and the cost to reach it. Routers that are trying to reach an external network use these advertisements to determine the best path to the next hop. ABRs generate Type 4 LSAs.
- **Autonomous system external LSA (Type 5)**—Redistributes routes from another autonomous system, usually from a different routing protocol into OSPF. In OSPFv3, addresses for these LSAs are expressed as “prefix and prefix length” instead of “address and mask.” The default route is expressed as a prefix with length 0.
- **Autonomous system external LSA (Type 7)**—Provides for carrying external route information within an NSSA. Type 7 LSAs may be originated by and advertised throughout an NSSA. NSSAs do not receive or originate Type 5 LSAs. Type 7 LSAs are advertised only within a single NSSA. They are not flooded into the backbone area or into any other area by border routers.
- **Link LSA (Type 8)**—Has link-local flooding scope and is never flooded beyond the link with which it is associated. Link LSAs provide the link-local address of the router to all other routers attached to the link or network segment, inform other routers attached to the link of a list of IPv6 prefixes to associate with the link, and allow the router to assert a collection of Options bits to associate with the network LSA that is originated for the link.
- **Intra-area-prefix LSAs (Type 9)**—A router can originate multiple intra-area-prefix LSAs for every router or transit network, each with a unique link-state ID. The link-state ID for each intra-area-prefix LSA describes its association to either the router LSA or network LSA and contains prefixes for stub and transit networks.

An address prefix occurs in almost all newly defined LSAs. The prefix is represented by three fields: Prefix Length, Prefix Options, and Address Prefix. In OSPFv3, addresses for these LSAs are expressed as “prefix and prefix length” instead of “address and mask.” The default route is expressed as a prefix with length 0.

Inter-area-prefix and intra-area-prefix LSAs carry all IPv6 prefix information that, in IPv4, is included in router LSAs and network LSAs. The Options field in certain LSAs (router LSAs, network LSAs, interarea-router LSAs, and link LSAs) has been expanded to 24 bits to provide support for OSPF in IPv6.

In OSPFv3, the sole function of link-state ID in interarea-prefix LSAs, interarea-router LSAs, and autonomous system external LSAs is to identify individual pieces of the link-state database. All addresses or router IDs that are expressed by the link-state ID in OSPF Version 2 are carried in the body of the LSA in OSPFv3.

## Virtual Link and Transit Area for OSPF

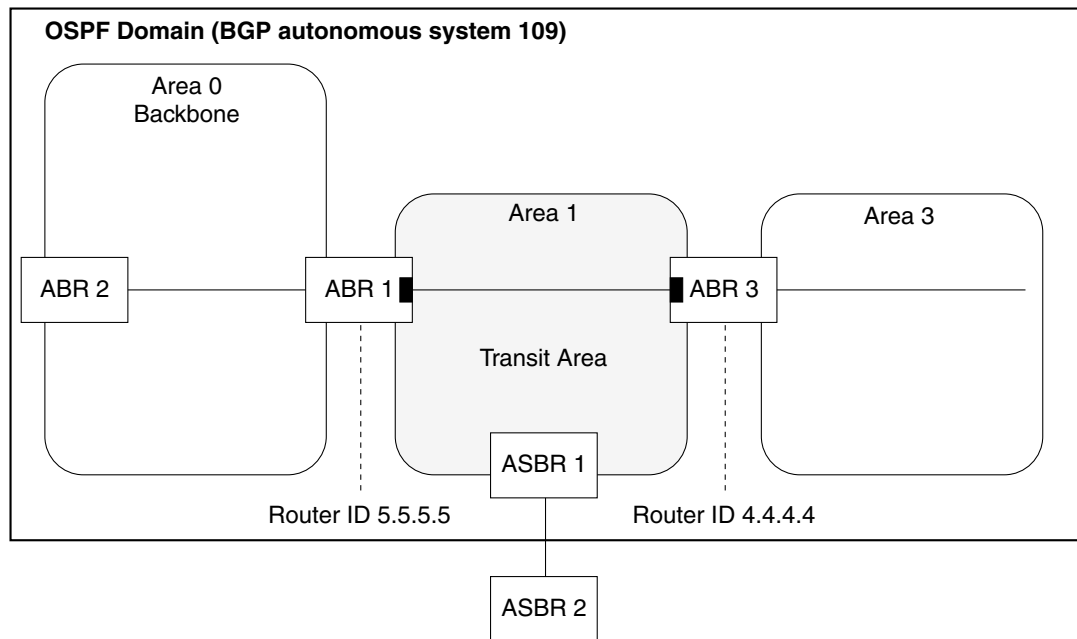
In OSPF, routing information from all areas is first summarized to the backbone area by ABRs. The same ABRs, in turn, propagate such received information to their attached areas. Such hierarchical distribution of routing information requires that all areas be connected to the backbone area (Area 0). Occasions might exist for which an area must be defined, but it cannot be physically connected to Area 0. Examples of such an occasion might be if your company makes a new acquisition that includes an OSPF area, or if Area 0 itself is partitioned.

In the case in which an area cannot be connected to Area 0, you must configure a virtual link between that area and Area 0. The two endpoints of a virtual link are ABRs, and the virtual link must be configured in both routers. The common nonbackbone area to which the two routers belong is called a transit area. A virtual link specifies the transit area and the router ID of the other virtual endpoint (the other ABR).

A virtual link cannot be configured through a stub area or NSSA.

Figure 2 illustrates a virtual link from Area 3 to Area 0.

**Figure 2** Virtual Link to Area 0



88722

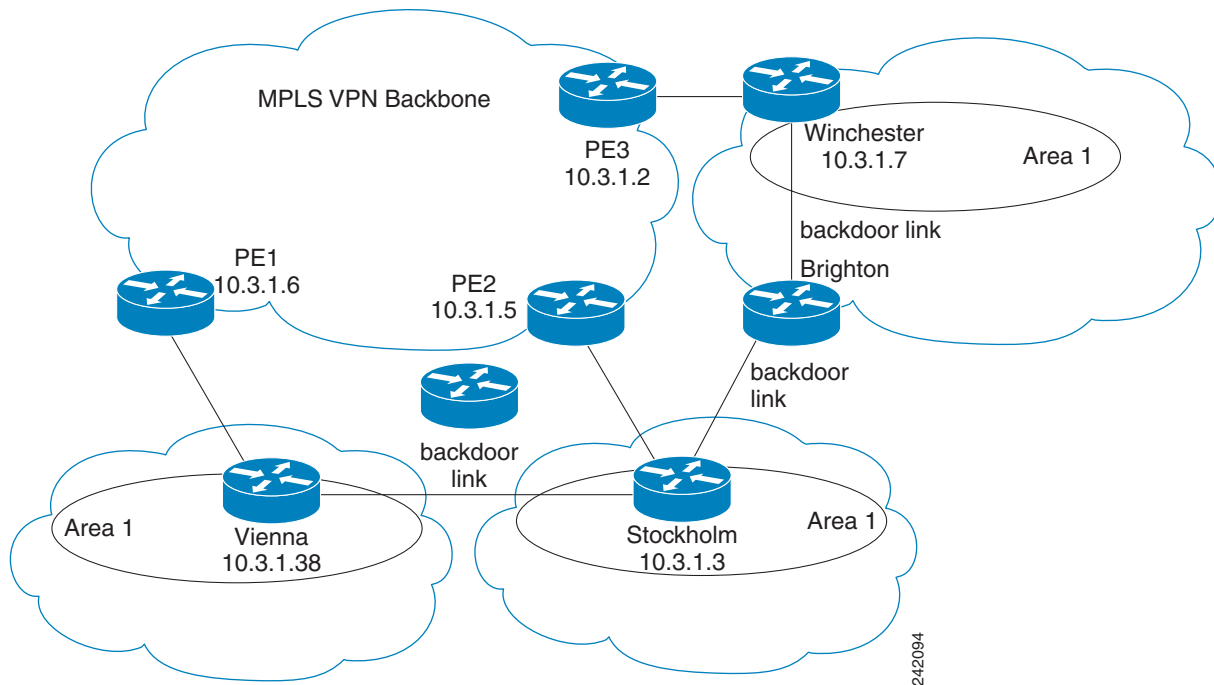
## OSPFv2 Shamlink Support for MPLS VPN

In an MPLS VPN environment, several VPN client sites can be connected in the same OSPF area. If these sites are connected over a backdoor link (intra-area link) and connected over the VPN backbone, all traffic passes over the backdoor link instead of over the VPN backbone, because provider edge routers advertise OSPF routes learned over the VPN backbone as inter-area or external routes that are less preferred than intra-area routes advertised over backdoor links.

To correct this default OSPF behavior in an MPLS VPN, configure a sham link between two provider edge (PE) routers to connect the sites through the MPLS VPN backbone. A sham link represents an intra-area (unnumbered point-to-point) connection between PE routers. All other routers in the area see the sham link and use it to calculate intra-area shortest path first (SPF) routes to the remote site. A cost must be configured with each sham link to determine whether traffic is sent over the backdoor link or sham link.

Configured source and destination addresses serve as the endpoints of the sham link. The source and destination IP addresses must belong to the VRF and be advertised by Border Gateway Protocol (BGP) as host routes to remote PE routers. The sham-link endpoint addresses should not be advertised by OSPF.

**Figure 3 Backdoor Paths Between OSPF Client Sites**



For example, [Figure 3](#) shows three client sites, each with backdoor links. Because each site runs OSPF within Area 1 configuration, all routing between the sites follows the intra-area path across the backdoor links instead of over the MPLS VPN backbone.

If the backdoor links between the sites are only used for backup purposes, default route selection over the backbone link is not acceptable as it creates undesirable traffic flow. To establish the desired path selection over the MPLS backbone, an additional OSPF intra-area (sham link) link between the ingress and egress PE routers must be created.

A sham link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link. If no backdoor link exists between sites, no sham link is required.

242094

**Figure 4** Sham Link Between PE Routers to Connected OSPF Client Sites

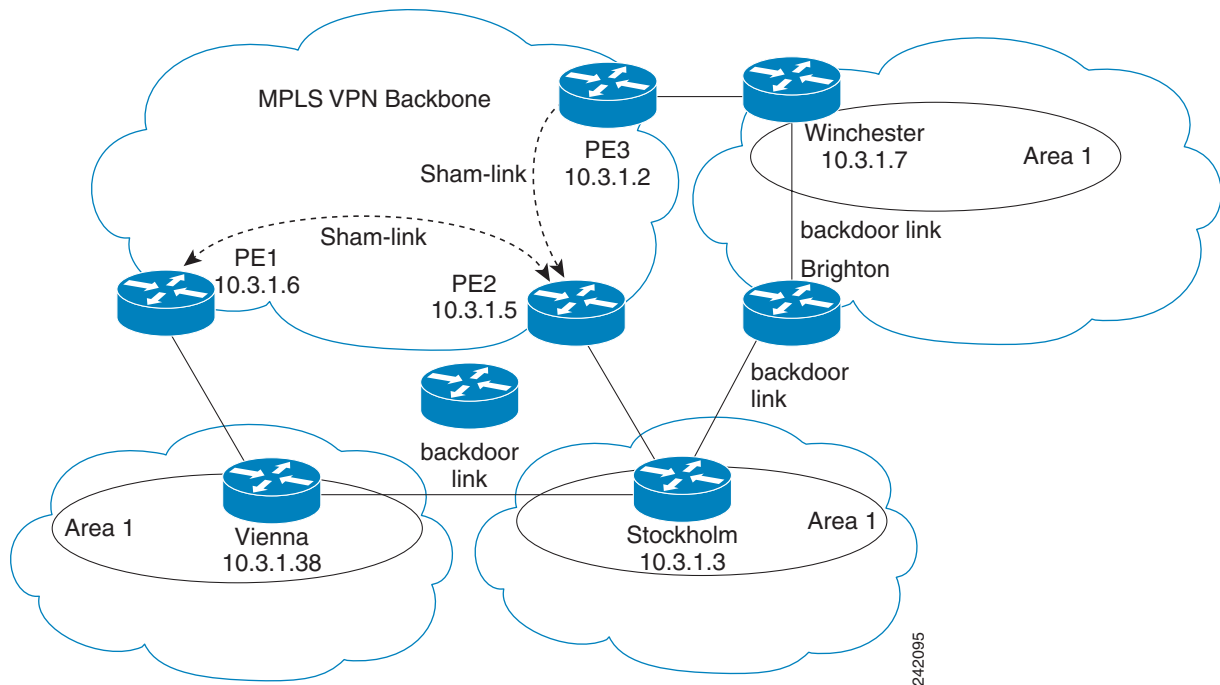


Figure 4 shows an MPLS VPN topology where a sham link configuration is necessary. A VPN client has three sites, each with a backdoor link. Two sham links are configured, one between PE-1 and PE-2 and another between PE-2 and PE-3. A sham link is not required between PE-1 and PE-3, because there is no backdoor link between these sites.

When a sham link is configured between the PE routers, the PE routers can populate the virtual routing and forwarding (VRF) table with the OSPF routes learned over the sham link. These OSPF routes have a larger administrative distance than BGP routes. If BGP routes are available, they are preferred over these OSPF routes with the high administrative distance.

## Route Redistribution for OSPF

Redistribution allows different routing protocols to exchange routing information. This technique can be used to allow connectivity to span multiple routing protocols. It is important to remember that the **redistribute** command controls redistribution *into* an OSPF process and not from OSPF. See the [“Configuration Examples for Implementing OSPF on Cisco IOS XR Software”](#) section on page RC-255 for an example of route redistribution for OSPF.

## OSPF Shortest Path First Throttling

OSPF SPF throttling makes it possible to configure SPF scheduling in millisecond intervals and to potentially delay SPF calculations during network instability. SPF is scheduled to calculate the Shortest Path Tree (SPT) when there is a change in topology. One SPF run may include multiple topology change events.



The interval at which the SPF calculations occur is chosen dynamically and based on the frequency of topology changes in the network. The chosen interval is within the boundary of the user-specified value ranges. If network topology is unstable, SPF throttling calculates SPF scheduling intervals to be longer until topology becomes stable.


SPF calculations occur at the interval set by the **timers throttle spf** command. The wait interval indicates the amount of time to wait until the next SPF calculation occurs. Each wait interval after that calculation is twice as long as the previous interval until the interval reaches the maximum wait time specified.

The SPF timing can be better explained using an example. In this example, the start interval is set at 5 milliseconds (ms), initial wait interval at 1000 ms, and maximum wait time at 90,000 ms.

```
timers spf 5 1000 90000
```

Figure 5 shows the intervals at which the SPF calculations occur as long as at least one topology change event is received in a given wait interval.

**Figure 5** *SPF Calculation Intervals Set by the timers spf Command*




Notice that the wait interval between SPF calculations doubles when at least one topology change event is received during the previous wait interval. After the maximum wait time is reached, the wait interval remains the same until the topology stabilizes and no event is received in that interval.

If the first topology change event is received after the current wait interval, the SPF calculation is delayed by the amount of time specified as the start interval. The subsequent wait intervals continue to follow the dynamic pattern.

If the first topology change event occurs after the maximum wait interval begins, the SPF calculation is again scheduled at the start interval and subsequent wait intervals are reset according to the parameters specified in the **timers throttle spf** command. Notice in Figure 6 that a topology change event was received after the start of the maximum wait time interval and that the SPF intervals have been reset.

**Figure 6** *Timer Intervals Reset After Topology Change Event*



## Nonstop Forwarding for OSPF Version 2

Cisco IOS XR NSF for OSPF Version 2 allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a failover. With NSF, peer networking devices do not experience routing flaps. During failover, data traffic is forwarded

through intelligent line cards while the standby Route Processor (RP) assumes control from the failed RP. The ability of line cards to remain up through a failover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to Cisco IOS XR NSF operation.

Routing protocols, such as OSPF, run only on the active RP or DRP and receive routing updates from their neighbor routers. When an OSPF NSF-capable router performs an RP failover, it must perform two tasks to resynchronize its link-state database with its OSPF neighbors. First, it must relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the link-state database for the network.

As quickly as possible after an RP failover, the NSF-capable router sends an OSPF NSF signal to neighboring NSF-aware devices. This signal is in the form of a link-local LSA generated by the failed-over router. Neighbor networking devices recognize this signal as a cue that the neighbor relationship with this router should not be reset. As the NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

After neighbor relationships are re-established, the NSF-capable router begins to resynchronize its database with all of its NSF-aware neighbors. At this point, the routing information is exchanged between the OSPF neighbors. After this exchange is completed, the NSF-capable device uses the routing information to remove stale routes, update the RIB, and update the FIB with the new forwarding information. OSPF on the router and the OSPF neighbors are now fully converged.

## Graceful Restart for OSPFv3

In the current release, various restart scenarios in the control plane of an IPv6-enabled router can disrupt data forwarding. The OSPFv3 Graceful Restart feature can preserve the data plane capability in the following circumstances:

- RP failure, resulting in a switchover to the backup processor
- Planned OSPFv3 process restart, such as software upgrade or downgrade
- Unplanned OSPFv3 process restart, such as a process crash

This feature supports nonstop data forwarding on established routes while the OSPFv3 routing protocol is restarting. (Therefore, this feature enhances high availability of IPv6 forwarding.)

## Modes of Graceful Restart Operation

The two operational modes that a router can be in for this feature are restart mode and helper mode. Restart mode occurs when the OSPFv3 process is doing a graceful restart. Helper mode refers to the neighbor routers that continue to forward traffic on established OSPFv3 routes while OSPFv3 is restarting on a neighboring router.

### Restart Mode

When the OSPFv3 process starts up, it determines whether it must attempt a graceful restart. The determination is based on whether graceful restart was previously enabled. (OSPFv3 does not attempt a graceful restart upon the first-time startup of the router.) When OSPFv3 graceful restart is enabled, it changes the purge timer in the RIB to a nonzero value. See [Configuring OSPFv3 Graceful Restart, page RC-230](#), for descriptions of how to enable and configure graceful restart.

During a graceful restart, the router does not populate OSPFv3 routes in the RIB. It tries to bring up full adjacencies with the fully adjacent neighbors that OSPFv3 had before the restart. Eventually, the OSPFv3 process indicates to the RIB that it has converged, either for the purpose of terminating the graceful restart (for any reason) or because it has completed the graceful restart.

The following are general details about restart mode. More detailed information on behavior and certain restrictions and requirements appears in the [Graceful Restart Requirements and Restrictions](#) section.

- If OSPFv3 attempts a restart too soon after the most recent restart, the OSPFv3 process is most likely crashing repeatedly, so the new graceful restart stops running. To control the period between allowable graceful restarts, use the **graceful-restart interval** command.
- When OSPFv3 starts a graceful restart with the first interface that comes up, a timer starts running to limit the duration (or lifetime) of the graceful restart. You can configure this period with the **graceful-restart lifetime** command. On each interface that comes up, a *grace* LSA (type 11) is flooded to indicate to the neighboring routers that this router is attempting graceful restart. The neighbors enter into helper mode.
- The designated router and backup designated router check of the hello packet received from the restarting neighbor is bypassed, because it might not be valid.

## Helper Mode

Helper mode is enabled by default. When a (helper) router receives a grace LSA (type 11) from a router that is attempting a graceful restart, the following events occur:

- If helper mode has been disabled through the **graceful-restart helper disable** command, the router drops the LSA packet.
- If helper mode is enabled, the router enters helper mode if all of the following conditions are met:
  - The local router itself is not attempting a graceful restart.
  - The local (helping) router has full adjacency with the sending neighbor.
  - The value of *lsage* (link state age) in the received LSA is less than the requested grace period.
  - The sender of the grace LSA is the same as the originator of the grace LSA.
- Upon entering helper mode, a router performs its helper function for a specific period of time. This time period is the lifetime value from the router that is in restart mode—minus the value of *lsage* in the received grace LSA. If the graceful restart succeeds in time, the helper's timer is stopped before it expires. If the helper's timer does expire, the adjacency to the restarting router is brought down, and normal OSPFv3 functionality resumes.
- The dead timer is not honored by the router that is in helper mode.
- A router in helper mode ceases to perform the helper function in any of the following cases:
  - The helper router is able to bring up a FULL adjacency with the restarting router.
  - The local timer for the helper function expires.

## Graceful Restart Requirements and Restrictions

The requirements for supporting the Graceful Restart feature include:

- Cooperation of a router's neighbors during a graceful restart. In relation to the router on which OSPFv3 is restarting, each router is called a *helper*.
- All neighbors of the router that does a graceful restart must be capable of doing a graceful restart.
- A graceful restart does not occur upon the first-time startup of a router.
- OSPFv3 neighbor information and database information are not check-pointed.
- An OSPFv3 process rebuilds adjacencies after it restarts.

- To ensure consistent databases after a restart, the OSPFv3 configuration must be identical to the configuration before the restart. (This requirement applies to self-originated information in the local database.) A graceful restart can fail if configurations change during the operation. In this case, data forwarding would be affected. OSPFv3 resumes operation by regenerating all its LSAs and resynchronizing its database with all its neighbors.
- Although IPv6 FIB tables remain unchanged during a graceful restart, these tables eventually mark the routes as stale through the use of a holddown timer. Enough time is allowed for the protocols to rebuild state information and converge.
- The router on which OSPFv3 is restarting must send OSPFv3 hellos within the dead interval of the process restart. Protocols must be able to retain adjacencies with neighbors before the adjacency dead timer expires. The default for the dead timer is 40 seconds. If hellos do not arrive on the adjacency before the dead timer expires, the router takes down the adjacency. The OSPFv3 Graceful Restart feature does not function properly if the dead timer is configured to be less than the time required to send hellos after the OSPFv3 process restarts.
- Simultaneous graceful restart sessions on multiple routers are not supported on a single network segment. If a router determines that multiple routers are in restart mode, it terminates any local graceful restart operation.
- This feature utilizes the available support for changing the purge time of existing OSPFv3 routes in the Routing Information Base (RIB). When graceful restart is enabled, the purge timer is set to 90 seconds by default. If graceful restart is disabled, the purge timer setting is 0.
- This feature has an associated *grace* LSA. This link-scope LSA is type 11.
- According to the RFC, the OSPFv3 process should flush all old, self-originated LSAs during a restart. With the Graceful Restart feature, however, the router delays this flushing of unknown self-originated LSAs during a graceful restart. OSPFv3 can learn new information and build new LSAs to replace the old LSAs. When the delay is over, all old LSAs are flushed.
- If graceful restart is enabled, the adjacency creation time of all the neighbors is saved in the system database (SysDB). The purpose for saving the creation time is so that OSPFv3 can use the original adjacency creation time to display the uptime for that neighbor after the restart.

## Warm Standby and Nonstop Routing for OSPF Version 2

OSPFv2 warm standby provides high availability across RP switchovers. With warm standby extensions, each process running on the active RP has a corresponding standby process started on the standby RP. A standby OSPF process can send and receive OSPF packets with no performance impact to the active OSPF process.

Nonstop routing (NSR) allows an RP failover, process restart, or in-service upgrade to be invisible to peer routers and ensures that there is minimal performance or processing impact. Routing protocol interactions between routers are not impacted by NSR. NSR is built on the warm standby extensions. NSR alleviates the requirement for Cisco NSF and IETF graceful restart protocol extensions.

See *Cisco IOS XR IP Addresses and Services Configuration Guide* for more information on NSR.

## Multicast-Intact Support for OSPF

The multicast-intact feature provides the ability to run multicast routing (PIM) when IGP shortcuts are configured and active on the router. Both OSPFv2 and IS-IS support the multicast-intact feature.

You can enable multicast-intact in the IGP when multicast routing protocols (PIM) are configured and IGP shortcuts are configured on the router. IGP shortcuts are MPLS tunnels that are exposed to IGP. The IGP routes IP traffic over these tunnels to destinations that are downstream from the egress router of the tunnel (from an SPF perspective). PIM cannot use IGP shortcuts for propagating PIM joins, because reverse path forwarding (RPF) cannot work across a unidirectional tunnel.

When you enable multicast-intact on an IGP, the IGP publishes a parallel or alternate set of equal-cost next hops for use by PIM. These next hops are called *mcast-intact* next hops. The mcast-intact next hops have the following attributes:

- They are guaranteed not to contain any IGP shortcuts.
- They are not used for unicast routing but are used only by PIM to look up an IPv4 next-hop to a PIM source.
- They are not published to the FIB.
- When multicast-intact is enabled on an IGP, all IPv4 destinations that were learned through link-state advertisements are published with a set equal-cost mcast-intact next hops to the RIB. This attribute applies even when the native next hops have no IGP shortcuts.

In OSPF, the max-paths (number of equal-cost next hops) limit is applied separately to the native and mcast-intact next hops. The number of equal cost mcast-intact next hops is the same as that configured for the native next hops.

## Load Balancing in OSPF Version 2 and OSPFv3

When a router learns multiple routes to a specific network by using multiple routing processes (or routing protocols), it installs the route with the lowest administrative distance in the routing table. Sometimes the router must select a route from among many learned by using the same routing process with the same administrative distance. In this case, the router chooses the path with the lowest cost (or metric) to the destination. Each routing process calculates its cost differently; the costs may need to be manipulated to achieve load balancing.

OSPF performs load balancing automatically. If OSPF finds that it can reach a destination through more than one interface and each path has the same cost, it installs each path in the routing table. The only restriction on the number of paths to the same destination is controlled by the **maximum-paths** (OSPF) command. The default number of maximum paths is 32 for Cisco CRS-1 routers and 16 for Cisco XR 12000 Series Routers. The range is from 1 to 32 for Cisco CRS-1 routers and 1 to 16 for Cisco XR 12000 Series Routers.

## Multi-Area Adjacency for OSPF Version 2

The multi-area adjacency feature for OSPFv2 allows a link to be configured on the primary interface in more than one area so that the link could be considered as an intra-area link in those areas and configured as a preference over more expensive paths.

This feature establishes a point-to-point unnumbered link in an OSPF area. A point-to-point link provides a topological path for that area, and the primary adjacency uses the link to advertise the link consistent with draft-ietf-ospf-multi-area-adj-06.

The following are multi-area interface attributes and limitations:

- Exists as a logical construct over an existing primary interface for OSPF; however, the neighbor state on the primary interface is independent of the multi-area interface.

- Establishes a neighbor relationship with the corresponding multi-area interface on the neighboring router. A mixture of multi-area and primary interfaces is not supported.
- Advertises an unnumbered point-to-point link in the router link state advertisement (LSA) for the corresponding area when the neighbor state is full.
- Created as a point-to-point network type and is not configurable.
- Created only on native point-to-point interfaces, such as Packet-over-SONET (PoS) or serial.
- Inherits the Bidirectional Forwarding Detection (BFD) characteristics from its primary interface. BFD is not configurable under a multi-area interface; however, it is configurable under the primary interface.

The multi-area interface inherits the interface characteristics from its primary interface, but some interface characteristics can be configured under the multi-area interface configuration mode as shown below:

```
RP/0/RP0/CPU0:router(config-ospf-ar)# multi-area-interface POS 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-mif)# ?

  authentication          Enable authentication
  authentication-key      Authentication password (key)
  commit                  Commit the configuration changes to running
  cost                     Interface cost
  database-filter         Filter OSPF LSA during synchronization and flooding
  dead-interval           Interval after which a neighbor is declared dead
  describe                Describe a command without taking real actions
  distribute-list         Filter networks in routing updates
  do                       Run an exec command
  exit                    Exit from this submode
  hello-interval          Time between HELLO packets
  message-digest-key      Message digest authentication password (key)
  mtu-ignore              Enable/Disable ignoring of MTU in DBD packets
  no                       Negate a command or set its defaults
  packet-size             Customize size of OSPF packets upto MTU
  pwd                     Commands used to reach current submode
  retransmit-interval     Time between retransmitting lost link state advertisements
  root                    Exit to the global configuration mode
  show                    Show contents of configuration
  transmit-delay          Estimated time needed to send link-state update packet
RP/0/RP0/CPU0:router(config-ospf-ar-mif)#
```

## Label Distribution Protocol IGP Auto-configuration for OSPF

Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) auto-configuration simplifies the procedure to enable LDP on a set of interfaces used by an IGP instance, such as OSPF. LDP IGP auto-configuration can be used on a large number of interfaces (for example, when LDP is used for transport in the core) and on multiple OSPF instances simultaneously.

This feature supports the IPv4 unicast address family for the default VPN routing and forwarding (VRF) instance.

LDP IGP auto-configuration can also be explicitly disabled on an individual interface basis under LDP using the **ldp igp auto-config disable** command. This allows LDP to receive all OSPF interfaces minus the ones explicitly disabled.

See *Cisco IOS XR Multiprotocol Label Switching Configuration Guide* for information on configuring LDP IGP auto-configuration.

## OSPF Authentication Message Digest Management

All OSPF routing protocol exchanges are authenticated and the method used can vary depending on how authentication is configured. When using cryptographic authentication, the OSPF routing protocol uses the Message Digest 5 (MD5) authentication algorithm to authenticate packets transmitted between neighbors in the network. For each OSPF protocol packet, a key is used to generate and verify a message digest that is appended to the end of the OSPF packet. The message digest is a one-way function of the OSPF protocol packet and the secret key. Each key is identified by the combination of interface used and the key identification. An interface may have multiple keys active at any time.

To manage the rollover of keys and enhance MD5 authentication for OSPF, you can configure a container of keys called a *keychain* with each key comprising the following attributes: generate/accept time, key identification, and authentication algorithm.

## GTSM TTL Security Mechanism for OSPF

OSPF is a link state protocol that requires networking devices to detect topological changes in the network, flood Link State Advertisement (LSA) updates to neighbors, and quickly converge on a new view of the topology. However, during the act of receiving LSAs from neighbors, network attacks can occur, because there are no checks that unicast or multicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the Generalized TTL Security Mechanism (GTSM), RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to when OSPF packets are originated and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

## Path Computation Element for OSPFv2

A PCE is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCE is accomplished when a PCE address and client is configured for MPLS-TE. PCE communicates its PCE address and capabilities to OSPF then OSPF packages this information in the PCE Discovery type-length-value (TLV) (Type 2) and re originates the RI LSA. OSPF also includes the Router Capabilities TLV (Type 1) in all its RI LSAs. The PCE Discovery TLV contains the PCE address sub-TLV (Type 1) and the Path Scope Sub-TLV (Type 2).

The PCE Address Sub-TLV specifies the IP address that must be used to reach the PCE. It should be a loop-back address that is always reachable, this TLV is mandatory, and must be present within the PCE Discovery TLV. The Path Scope Sub-TLV indicates the PCE path computation scopes, which refers to the PCE ability to compute or participate in the computation of intra-area, inter-area, inter-AS or inter-layer TE LSPs.

PCE extensions to OSPFv2 include support for the Router Information Link State Advertisement (RI LSA). OSPFv2 is extended to receive all area scopes (LSA Types 9, 10, and 11). However, OSPFv2 originates only area scope Type 10.

For detailed information for the Path Computation Element feature see the “Implementing MPLS Traffic Engineering on Cisco IOS XR software” module of the *Cisco IOS XR MPLS Configuration Guide*, Release 3.6, and the following IETF drafts:

- [draft-ietf-ospf-cap-09](#)
- [draft-ietf-pce-disco-proto-ospf-00](#)

## How to Implement OSPF on Cisco IOS XR Software

This section contains the following procedures:

- [Enabling OSPF, page RC-194](#) (required)
- [Configuring Stub and Not-So-Stubby Area Types, page RC-197](#) (optional)
- [Configuring Neighbors for Nonbroadcast Networks, page RC-199](#) (optional)
- [Configuring Authentication at Different Hierarchical Levels for OSPF Version 2, page RC-204](#) (optional)
- [Controlling the Frequency That the Same LSA Is Originated or Accepted for OSPF, page RC-207](#) (optional)
- [Creating a Virtual Link with MD5 Authentication to Area 0 for OSPF, page RC-209](#) (optional)
- [Summarizing Subnetwork LSAs on an OSPF ABR, page RC-213](#) (optional)
- [Redistributing Routes from One IGP into OSPF, page RC-215](#) (optional)
- [Configuring OSPF Shortest Path First Throttling, page RC-219](#) (optional)
- [Configuring Nonstop Forwarding Specific to Cisco for OSPF Version 2, page RC-224](#) (optional)
- [Configuring OSPF Version 2 for MPLS Traffic Engineering, page RC-226](#) (optional)
- [Configuring OSPFv3 Graceful Restart, page RC-230](#) (optional)
- [Configuring an OSPFv2 Sham Link, page RC-233](#) (optional)
- [Enabling Nonstop Routing for OSPFv2, page RC-236](#) (optional)
- [Enabling Multicast-intact for OSPFv2, page RC-237](#) (optional)
- [Associating Interfaces to a VRF, page RC-238](#) (optional)
- [Configuring OSPF as a Provider Edge to Customer Edge \(PE-CE\) Protocol, page RC-240](#) (optional)
- [Creating Multiple OSPF Instances \(OSPF Process and a VRF\), page RC-243](#) (optional)
- [Configuring Multi-area Adjacency, page RC-244](#) (optional)
- [Configuring Label Distribution Protocol IGP Auto-configuration for OSPF, page RC-246](#) (optional)
- [Configuring Authentication Message Digest Management for OSPF, page RC-247](#) (optional)
- [Configuring Generalized TTL Security Mechanism \(GTSM\) for OSPF, page RC-251](#) (optional)
- [Verifying OSPF Configuration and Operation, page RC-254](#) (optional)

### Enabling OSPF

This task explains how to perform the minimum OSPF configuration on your router that is to enable an OSPF process with a router ID, configure a backbone or nonbackbone area, and then assign one or more interfaces on which OSPF runs.



## Prerequisites

Although you can configure OSPF before you configure an IP address, no OSPF routing occurs until at least one IP address is configured.

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*  
or  
**router ospfv3** *process-name*
3. **router-id** {*router-id*}
4. **area** *area-id*
5. **interface** *type instance*
6. Repeat Step 5 for each interface that uses OSPF.
7. **log adjacency changes** [**detail**] [**enable** | **disable**]
8. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  or Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IP address as the router ID.
	<b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	

	Command or Action	Purpose
Step 4	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures an area for the OSPF process.</p> <ul style="list-style-type: none"> <li>• Backbone areas have an area ID of 0.</li> <li>• Nonbackbone areas have a nonzero area ID.</li> <li>• The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 5	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3</p>	<p>Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.</p>
Step 6	Repeat Step 5 for each interface that uses OSPF.	—
Step 7	<p><b>log adjacency changes</b> [<b>detail</b>] [<b>enable</b>   <b>disable</b>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# log adjacency changes detail</p>	<p>(Optional) Requests notification of neighbor changes.</p> <ul style="list-style-type: none"> <li>• By default, this feature is enabled.</li> <li>• The messages generated by neighbor changes are considered notifications, which are categorized as severity Level 5 in the <b>logging console</b> command. The <b>logging console</b> command controls which severity level of messages are sent to the console. By default, all severity level messages are sent.</li> </ul>
Step 8	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Stub and Not-So-Stubby Area Types

This task explains how to configure the stub area and the NSSA for OSPF.

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*  
or  
**router ospfv3** *process-name*
3. **router-id** {*router-id*}
4. **area** *area-id*
5. **stub** [no-summary]  
or  
**nssa** [no-redistribution] [default-information-originate] [no-summary]
6. **stub**  
or  
**nssa**
7. **default-cost** *cost*
8. **end**  
or  
**commit**
9. Repeat this task on all other routers in the stub area or NSSA.

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  or Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode. <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IP address as the router ID.

	Command or Action	Purpose
Step 4	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1</p>	<p>Enters area configuration mode and configures a nonbackbone area for the OSPF process.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 5	<p><b>stub</b> [<b>no-summary</b>] OR <b>nssa</b> [<b>no-redistribution</b>] [<b>default-information-originate</b>] [<b>no-summary</b>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# stub no summary OR RP/0/RP0/CPU0:router(config-ospf-ar)# nssa no-redistribution</p>	<p>Defines the nonbackbone area as a stub area.</p> <ul style="list-style-type: none"> <li>See the “<a href="#">Configuring Stub and Not-So-Stubby Area Types</a>” section on page RC-197.</li> <li>Specify the <b>no-summary</b> keyword to further reduce the number of LSAs sent into a stub area. This keyword prevents the ABR from sending summary link-state advertisements (Type 3) in the stub area.</li> </ul> <p>or</p> <p>Defines an area as an NSSA.</p> <ul style="list-style-type: none"> <li>See the “<a href="#">Configuring Stub and Not-So-Stubby Area Types</a>” section on page RC-197.</li> </ul>
Step 6	<p><b>stub</b> OR <b>nssa</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# stub OR RP/0/RP0/CPU0:router(config-ospf-ar)# nssa</p>	<p>(Optional) Turns off the options configured for stub and NSSA areas.</p> <ul style="list-style-type: none"> <li>If you configured the stub and NSSA areas using the optional keywords (<b>no-summary</b>, <b>no-redistribution</b>, <b>default-information-originate</b>, and <b>no-summary</b>) in Step 5, you must now reissue the <b>stub</b> and <b>nssa</b> commands without the keywords—rather than using the <b>no</b> form of the command.</li> <li>For example, the <b>no nssa default-information-originate</b> form of the command changes the NSSA area into a normal area that inadvertently brings down the existing adjacencies in that area.</li> </ul>
Step 7	<p><b>default-cost</b> <i>cost</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# default-cost 15</p>	<p>(Optional) Specifies a cost for the default summary route sent into a stub area or an NSSA.</p> <ul style="list-style-type: none"> <li>Use this command only on ABRs attached to the NSSA. Do not use it on any other routers in the area.</li> <li>The default cost is 1.</li> </ul>

	Command or Action	Purpose
Step 8	<pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# end or RP/0/RP0/CPU0:router(config-ospf-ar)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 9	Repeat this task on all other routers in the stub area or NSSA.	—

## Configuring Neighbors for Nonbroadcast Networks

This task explains how to configure neighbors for a nonbroadcast network. This task is optional.

### Prerequisites

Configuring NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network.

### SUMMARY STEPS

- configure**
- router ospf** *process-name*  
or  
**router ospfv3** *process-name*
- router-id** {*router-id*}
- area** *area-id*
- network** {**broadcast** | **non-broadcast** | {**point-to-multipoint** [**non-broadcast**] | **point-to-point**}}
- dead-interval** *seconds*
- hello-interval** *seconds*
- interface** *type instance*

9. **neighbor** *ip-address* [**priority number**] [**poll-interval seconds**] [**cost number**]  
or  
**neighbor** *ipv6-link-local-address* [**priority number**] [**poll-interval seconds**] [**cost number**]  
[**database-filter** [all]]
10. Repeat Step 9 for all neighbors on the interface.
11. **exit**
12. **interface** *type instance*
13. **neighbor** *ip-address* [**priority number**] [**poll-interval seconds**][**cost number**] [**database-filter** [all]]  
or  
**neighbor** *ipv6-link-local-address* [**priority number**] [**poll-interval seconds**][**cost number**]  
[**database-filter** [all]]
14. Repeat Step 13 for all neighbors on the interface.
15. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  or Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode. <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	Configures a router ID for the OSPF process. <b>Note</b> We recommend using a stable IP address as the router ID.
Step 4	<b>area</b> <i>area-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode and configures an area for the OSPF process. <ul style="list-style-type: none"><li>• The example configures a backbone area.</li><li>• The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li></ul>

	Command or Action	Purpose
Step 5	<p><b>network</b> {<b>broadcast</b>   <b>non-broadcast</b>   {<b>point-to-multipoint</b> [<b>non-broadcast</b>]   <b>point-to-point</b>}}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# network non-broadcast</p>	<p>Configures the OSPF network type to a type other than the default for a given medium.</p> <ul style="list-style-type: none"> <li>The example sets the network type to NBMA.</li> </ul>
Step 6	<p><b>dead-interval</b> <i>seconds</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# dead-interval 40</p>	<p>(Optional) Sets the time to wait for a hello packet from a neighbor before declaring the neighbor down.</p>
Step 7	<p><b>hello-interval</b> <i>seconds</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# hello-interval 10</p>	<p>(Optional) Specifies the interval between hello packets that OSPF sends on the interface.</p>
Step 8	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/2/0/0</p>	<p>Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.</p> <ul style="list-style-type: none"> <li>In this example, the interface inherits the nonbroadcast network type and the hello and dead intervals from the areas because the values are not set at the interface level.</li> </ul>

Command or Action	Purpose
<p><b>Step 9</b></p> <pre>neighbor ip-address [priority number] [poll-interval seconds][cost number] or neighbor ipv6-link-local-address [priority number] [poll-interval seconds][cost number] [database-filter [all]]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-if)#  neighbor 10.20.20.1 priority 3 poll-interval 15  or  RP/0/RP0/CPU0:router(config-ospf-ar-if)#  neighbor fe80::3203:a0ff:fe9d:f3fe</p>	<p>Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.</p> <p>or</p> <p>Configures the link-local IPv6 address of OSPFv3 neighbors.</p> <ul style="list-style-type: none"> <li>The <i>ipv6-link-local-address</i> argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.</li> <li>The <b>priority</b> keyword notifies the router that this neighbor is eligible to become a DR or BDR. The priority value should match the actual priority setting on the neighbor router. The neighbor priority default value is zero. This keyword does not apply to point-to-multipoint interfaces.</li> <li>The <b>poll-interval</b> keyword does not apply to point-to-multipoint interfaces. RFC 1247 recommends that this value be much larger than the hello interval. The default is 120 seconds (2 minutes).</li> <li>Neighbors with no specific cost configured assumes the cost of the interface, based on the <b>cost</b> command. On point-to-multipoint interfaces, <b>cost number</b> is the only keyword and argument combination that works. The <b>cost</b> keyword does not apply to NBMA networks.</li> <li>The <b>database-filter</b> keyword filters outgoing LSAs to an OSPF neighbor. If you specify the <b>all</b> keyword, incoming and outgoing LSAs are filtered. Use with extreme caution since filtering may cause the routing topology to be seen as entirely different between two neighbors, resulting in ‘black-holing’ of data traffic or routing loops.</li> </ul>
<p><b>Step 10</b></p> <p>Repeat Step 9 for all neighbors on the interface.</p>	<p>—</p>
<p><b>Step 11</b></p> <pre>exit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit</p>	<p>Enters area configuration mode.</p>
<p><b>Step 12</b></p> <pre>interface type instance</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar)# interface  POS 0/3/0/1</p>	<p>Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.</p> <ul style="list-style-type: none"> <li>In this example, the interface inherits the nonbroadcast network type and the hello and dead intervals from the areas because the values are not set at the interface level.</li> </ul>



Command or Action	Purpose
<p><b>Step 13</b> <code>neighbor ip-address [priority number]</code>  <code>[poll-interval seconds] [cost number]</code>  <code>[database-filter [all]]</code>  or  <code>neighbor ipv6-link-local-address [priority</code>  <code>number] [poll-interval seconds] [cost number]</code>  <code>[database-filter [all]]</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar)# neighbor  10.34.16.6  or  RP/0/RP0/CPU0:router(config-ospf-ar)# neighbor  fe80::3203:a0ff:fe9d:f3f</p>	<p>Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.</p> <p>or</p> <p>Configures the link-local IPv6 address of OSPFv3 neighbors.</p> <ul style="list-style-type: none"> <li>• The <code>ipv6-link-local-address</code> argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.</li> <li>• The <b>priority</b> keyword notifies the router that this neighbor is eligible to become a DR or BDR. The priority value should match the actual priority setting on the neighbor router. The neighbor priority default value is zero. This keyword does not apply to point-to-multipoint interfaces.</li> <li>• The <b>poll-interval</b> keyword does not apply to point-to-multipoint interfaces. RFC 1247 recommends that this value be much larger than the hello interval. The default is 120 seconds (2 minutes).</li> <li>• Neighbors with no specific cost configured assumes the cost of the interface, based on the <b>cost</b> command. On point-to-multipoint interfaces, <b>cost number</b> is the only keyword and argument combination that works. The <b>cost</b> keyword does not apply to NBMA networks.</li> <li>• The <b>database-filter</b> keyword filters outgoing LSAs to an OSPF neighbor. If you specify the <b>all</b> keyword, incoming and outgoing LSAs are filtered. Use with extreme caution since filtering may cause the routing topology to be seen as entirely different between two neighbors, resulting in ‘black-holing’ or routing loops.</li> </ul>

Command or Action	Purpose
Step 14 Repeat Step 13 for all neighbors on the interface.	—
<p>Step 15</p> <pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# end or RP/0/RP0/CPU0:router(config-ospf-ar)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Authentication at Different Hierarchical Levels for OSPF Version 2

This task explains how to configure MD5 (secure) authentication on the OSPF router process, configure one area with plain text authentication, and then apply one interface with clear text (null) authentication.



### Note

Authentication configured at the interface level overrides authentication configured at the area level and the router process level. If an interface does not have authentication specifically configured, the interface inherits the authentication parameter value from a higher hierarchical level. See the [“OSPF Hierarchical CLI and CLI Inheritance”](#) section on page RC-176 for more information about hierarchy and inheritance.

## Prerequisites

If you choose to configure authentication, you must first decide whether to configure plain text or MD5 authentication, and whether the authentication applies to all interfaces in a process, an entire area, or specific interfaces. See the [“Route Authentication Methods for OSPF”](#) section on page RC-180 for information about each type of authentication and when you should use a specific method for your network.

## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **router-id** {*router-id*}
4. **authentication** [message-digest | null]

5. **message-digest-key** *key-id* **md5** {*key* | **clear** *key* | **encrypted** *key*}
6. **area** *area-id*
7. **interface** *type instance*
8. Repeat Step 7 for each interface that must communicate, using the same authentication.
9. **exit**
10. **area** *area-id*
11. **authentication** [**message-digest** | **null**]
12. **interface** *type instance*
13. Repeat Step 12 for each interface that must communicate, using the same authentication.
14. **interface** *type instance*
15. **authentication** [**message-digest** | **null**]
16. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	Configures a router ID for the OSPF process.
Step 4	<b>authentication</b> [ <b>message-digest</b>   <b>keychain</b> <i>keychain</i> ]   <b>null</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# authentication message-digest	Enables MD5 authentication for the OSPF process.  <ul style="list-style-type: none"> <li>This authentication type applies to the entire router process unless overridden by a lower hierarchical level such as the area or interface.</li> </ul>
Step 5	<b>message-digest-key</b> <i>key-id</i> <b>md5</b> { <i>key</i>   <b>clear</b> <i>key</i>   <b>encrypted</b> <i>key</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# message-digest-key 4 md5 yourkey	Specifies the MD5 authentication key for the OSPF process.  <ul style="list-style-type: none"> <li>The neighbor routers must have the same key identifier.</li> </ul>

	Command or Action	Purpose
Step 6	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	Enters area configuration mode and configures a backbone area for the OSPF process.
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3</p>	<p>Enters interface configuration mode and associates one or more interfaces to the backbone area.</p> <ul style="list-style-type: none"> <li>All interfaces inherit the authentication parameter values specified for the OSPF process (Step 4, Step 5, and Step 6).</li> </ul>
Step 8	Repeat Step 7 for each interface that must communicate, using the same authentication.	—
Step 9	<p><b>exit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# exit</p>	Enters area OSPF configuration mode.
Step 10	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1</p>	<p>Enters area configuration mode and configures a nonbackbone area 1 for the OSPF process.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 11	<p><b>authentication</b> [<b>message-digest</b>   <b>null</b>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# authentication</p>	<p>Enables Type 1 (plain text) authentication that provides no security.</p> <ul style="list-style-type: none"> <li>The example specifies plain text authentication (by not specifying a keyword). Use the <b>authentication-key</b> command in interface configuration mode to specify the plain text password.</li> </ul>
Step 12	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/0</p>	<p>Enters interface configuration mode and associates one or more interfaces to the nonbackbone area 1 specified in Step 7.</p> <ul style="list-style-type: none"> <li>All interfaces configured inherit the authentication parameter values configured for area 1.</li> </ul>
Step 13	Repeat Step 12 for each interface that must communicate, using the same authentication.	—
Step 14	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/3/0/0</p>	Enters interface configuration mode and associates one or more interfaces to a different authentication type.

	Command or Action	Purpose
Step 15	<p><b>authentication</b> [<b>message-digest</b>   <b>keychain</b> <i>keychain</i>]   <b>null</b>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-if)#  authentication null</p>	<p>Specifies no authentication on POS interface 0/3/0/0, overriding the plain text authentication specified for area 1.</p> <ul style="list-style-type: none"> <li>By default, all of the interfaces configured in the same area inherit the same authentication parameter values of the area.</li> </ul>
Step 16	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-if)# end  or  RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Controlling the Frequency That the Same LSA Is Originated or Accepted for OSPF

This task explains how to tune the convergence time of OSPF routes in the routing table when many LSAs need to be flooded in a very short time interval.

### SUMMARY STEPS

- configure**
- router ospf** *process-name*  
or  
**router ospfv3** *process-name*
- router-id** {*router-id*}
- Perform Step 5 or Step 6 or both to control the frequency that the same LSA is originated or accepted.
- timers lsa gen-interval** *seconds*
- timers lsa min-arrival** *seconds*
- timers lsa group-pacing** *seconds*

```

8. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  or Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IP address as the router ID.
Step 4	Perform Step 5 or Step 6 or both to control the frequency that the same LSA is originated or accepted.	—
Step 5	<b>timers lsa gen-interval</b> <i>seconds</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# timers lsa gen-interval 10	Changes the minimum interval between the same OSPF LSAs that the router originates. <ul style="list-style-type: none"><li>The default is 5 seconds for both OSPF and OSPFv3.</li></ul>
Step 6	<b>timers lsa min-arrival</b> <i>seconds</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# timers lsa min-arrival 2	Limits the frequency that new processes of any particular OSPF Version 2 LSA can be accepted during flooding. <ul style="list-style-type: none"><li>The default is 1 second.</li></ul>

	Command or Action	Purpose
Step 7	<p><code>timers lsa group-pacing seconds</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf)# timers lsa group-pacing 1000</p>	<p>Changes the interval at which OSPF link-state LSAs are collected into a group for flooding.</p> <ul style="list-style-type: none"> <li>The default is 240 seconds.</li> </ul>
Step 8	<p><code>end</code>  OR  <code>commit</code></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf)# end  OR  RP/0/RP0/CPU0:router(config-ospf)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Creating a Virtual Link with MD5 Authentication to Area 0 for OSPF

This task explains how to create a virtual link to your backbone (area 0) and apply MD5 authentication. You must perform the steps described on both ABRs, one at each end of the virtual link. To understand virtual links, see the [“Virtual Link and Transit Area for OSPF”](#) section on page RC-183.



### Note

After you explicitly configure area parameter values, they are inherited by all interfaces bound to that area—unless you override the values and configure them explicitly for the interface. An example is provided in the [“Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example”](#) section on page RC-260.

## Prerequisites

The following prerequisites must be met before creating a virtual link with MD5 authentication to area 0:

- You must have the router ID of the neighbor router at the opposite end of the link to configure the local router. You can execute the `show ospf` or `show ospfv3` command on the remote router to get its router ID.

- For a virtual link to be successful, you need a stable router ID at each end of the virtual link. You do not want them to be subject to change, which could happen if they are assigned by default (See the “[OSPF Process and Router ID](#)” section on page RC-179 for an explanation of how the router ID is determined.) Therefore, we recommend that you perform one of the following tasks before configuring a virtual link:
  - Use the **router-id** command to set the router ID. This strategy is preferable.
  - Configure a loopback interface so that the router has a stable router ID.
- Before configuring your virtual link for OSPF Version 2, you must decide whether to configure plain text authentication, MD5 authentication, or no authentication (which is the default). Your decision determines whether you need to perform additional tasks related to authentication.

**Note**


---

If you decide to configure plain text authentication or no authentication, see the **authentication** command provided in the *OSPF Commands on Cisco IOS XR Software* module in the *Cisco IOS XR Routing Command Reference*.

---

**SUMMARY STEPS**

1. **show ospf** *[process-name]*  
or  
**show ospfv3** *[process-name]*
2. **configure**
3. **router ospf** *process-name*  
or  
**router ospfv3** *process-name*
4. **router-id** {*router-id*}
5. **area** *area-id*
6. **virtual link** *router-id*
7. **authentication message-digest**
8. **message-digest-key** *key-id md5* {*key* | **clear** *key* | **encrypted** *key*}
9. Repeat all the steps in this task on the ABR that is at the other end of the virtual link. Specify the same key ID and key that you specified for the virtual link on this router.
10. **end**  
or  
**commit**
11. **show ospf** *[process-name]* [*area-id*] **virtual-links**  
or  
**show ospfv3** *[process-name]* **virtual-links**



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show ospf [process-name] OR show ospfv3 [process-name]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show ospf OR RP/0/RP0/CPU0:router# show ospfv3 </p>	<p>(Optional) Displays general information about OSPF routing processes.</p> <ul style="list-style-type: none"> <li>The output displays the router ID of the local router. You need this router ID to configure the other end of the link.</li> </ul>
Step 2	<pre>configure</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure </p>	Enters global configuration mode.
Step 3	<pre>router ospf process-name OR router ospfv3 process-name</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 OR RP/0/RP0/CPU0:router(config)# router ospfv3 1 </p>	<p>Enables OSPF routing for the specified routing process and places the router in router configuration mode.</p> <p>or</p> <p>Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.</p> <p><b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.</p>
Step 4	<pre>router-id {router-id}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3 </p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 5	<pre>area area-id</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1 </p>	<p>Enters area configuration mode and configures a nonbackbone area for the OSPF process.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 6	<pre>virtual-link router-id</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# virtual link 10.3.4.5 </p>	<p>Defines an OSPF virtual link.</p> <ul style="list-style-type: none"> <li>See the <a href="#">“Virtual Link and Transit Area for OSPF” section on page RC-183</a>.</li> </ul>
Step 7	<pre>authentication message-digest</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-vl)# authentication message-digest </p>	Selects MD5 authentication for this virtual link.

	Command or Action	Purpose
Step 8	<pre>message-digest-key key-id md5 {key   clear key   encrypted key}</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-v1)#  message-digest-key 4 md5 yourkey</p>	<p>Defines an OSPF virtual link.</p> <ul style="list-style-type: none"> <li>See the “<a href="#">Virtual Link and Transit Area for OSPF</a>” section on page RC-183 to understand a virtual link.</li> <li>The <i>key-id</i> argument is a number in the range from 1 to 255. The <i>key</i> argument is an alphanumeric string of up to 16 characters. The routers at both ends of the virtual link must have the same key identifier and key to be able to route OSPF traffic.</li> <li>The <b>authentication-key</b> <i>key</i> command is not supported for OSPFv3.</li> <li>Once the key is encrypted it must remain encrypted.</li> </ul>
Step 9	Repeat all of the steps in this task on the ABR that is at the other end of the virtual link. Specify the same key ID and key that you specified for the virtual link on this router.	—
Step 10	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-v1)# end  or  RP/0/RP0/CPU0:router(config-ospf-ar-v1)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 11	<pre>show ospf [process-name] [area-id] virtual-links or show ospfv3 [process-name] virtual-links</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show ospf 1 2  virtual-links  or  RP/0/RP0/CPU0:router# show ospfv3 1  virtual-links</p>	(Optional) Displays the parameters and the current state of OSPF virtual links.

## Examples

In the following example, the **show ospfv3 virtual links EXEC** command verifies that the OSPF\_VL0 virtual link to the OSPFv3 neighbor is up, the ID of the virtual link interface is 2, and the IPv6 address of the virtual link endpoint is 2003:3000::1.

```
RP/0/RP0/CPU0:router# show ospfv3 virtual-links

Virtual Links for OSPFv3 1

Virtual Link OSPF_VL0 to router 10.0.0.3 is up
  Interface ID 2, IPv6 address 2003:3000::1
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 0.1.20.255, via interface POS 0/1/0/1, Cost of using 2
  Transmit Delay is 5 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Adjacency State FULL (Hello suppressed)
  Index 0/2/3, retransmission queue length 0, number of retransmission 1
  First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

Check for lines:
Virtual Link OSPF_VL0 to router 10.0.0.3 is up
  Adjacency State FULL (Hello suppressed)

State is up and Adjacency State is FULL
```

## Summarizing Subnetwork LSAs on an OSPF ABR

If you configured two or more subnetworks when you assigned your IP addresses to your interfaces, you might want the software to summarize (aggregate) into a single LSA all of the subnetworks that the local area advertises to another area. Such summarization would reduce the number of LSAs and thereby conserve network resources. This summarization is known as interarea route summarization. It applies to routes from within the autonomous system. It does not apply to external routes injected into OSPF by way of redistribution.

This task configures OSPF to summarize subnetworks into one LSA, by specifying that all subnetworks that fall into a range are advertised together. This task is performed on an ABR only.

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*  
or  
**router ospfv3** *process-name*
3. **router-id** {*router-id*}
4. **area** *area-id*
5. **range** *ip-address mask* [**advertise** | **not-advertise**]  
or  
**range** *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]
6. **interface** *type instance*

```

7. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>configure</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<pre>router ospf process-name</pre> <p>or</p> <pre>router ospfv3 process-name</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1</p>	<p>Enables OSPF routing for the specified routing process and places the router in router configuration mode.</p> <p>or</p> <p>Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.</p> <p><b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.</p>
Step 3	<pre>router-id {router-id}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 4	<pre>area area-id</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures a nonbackbone area for the OSPF process.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 5	<pre>range ip-address mask [advertise   not-advertise]</pre> <p>or</p> <pre>range ipv6-prefix/prefix-length [advertise   not-advertise]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# range 192.168.0.0 255.255.0.0 advertise or</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# range 4004:f000::/32 advertise</p>	<p>Consolidates and summarizes OSPF routes at an area boundary.</p> <ul style="list-style-type: none"> <li>The <b>advertise</b> keyword causes the software to advertise the address range of subnetworks in a Type 3 summary LSA.</li> <li>The <b>not-advertise</b> keyword causes the software to suppress the Type 3 summary LSA, and the subnetworks in the range remain hidden from other areas.</li> <li>In the first example, all subnetworks for network 192.168.0.0 are summarized and advertised by the ABR into areas outside the backbone.</li> <li>In the second example, two or more IPv4 interfaces are covered by a 192.x.x network.</li> </ul>

	Command or Action	Purpose
Step 6	<pre><b>interface</b> <i>type instance</i></pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/2/0/3</p>	Enters interface configuration mode and associates one or more interfaces to the area.
Step 7	<pre><b>end</b> or <b>commit</b></pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# end or RP/0/RP0/CPU0:router(config-ospf-ar)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing Routes from One IGP into OSPF

This task redistributes routes from an IGP (could be a different OSPF process) into OSPF.

### Prerequisites

For information about configuring routing policy, see the *Implementing Routing Policy on Cisco IOS XR Software* module.

### SUMMARY STEPS

- configure**
- router ospf** *process-name*  
or  
**router ospfv3** *process-name*
- router-id** {*router-id*}
- redistribute** *protocol* [*process-id*] {**level-1** | **level-1-2** | **level-2**} [**metric** *metric-value*] [**metric-type** *type-value*] [**match** {**internal** | **external** [1 | 2] | **nssa-external** [1 | 2]}] [**tag** *tag-value*] [**route-map** *map-tag* | **route-policy** *policy-tag*]

5. **summary-prefix** *address mask* [**not-advertise**] [**tag** *tag*]  
or  
**summary-prefix** *ipv6-prefix/prefix-length* [**not-advertise**] [**tag** *tag*]
6. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode. or Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode. <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	Configures a router ID for the OSPF process. <b>Note</b> We recommend using a stable IPv4 address as the router ID.

Command or Action	Purpose
<p><b>Step 4</b></p> <pre> <b>redistribute</b> <i>protocol</i> [<i>process-id</i>] {<b>level-1</b>   <b>level-1-2</b>   <b>level-2</b>} [<b>metric</b> <i>metric-value</i>] [<b>metric-type</b> <i>type-value</i>] [<b>match</b> {<b>internal</b>   <b>external</b> [1   2]}   <b>nssa-external</b> [1   2]] [<b>tag</b> <i>tag-value</i>] [<b>route-map</b> <i>map-tag</i>   <b>policy</b> <i>policy-tag</i>]  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# redistribute bgp 1 level-1 or RP/0/RP0/CPU0:router(config-router)# redistribute bgp 1 level-1-2 metric-type 1                     </pre>	<p>Redistributes OSPF routes from one routing domain to another routing domain.</p> <p>or</p> <p>Redistributes OSPFv3 routes from one routing domain to another routing domain.</p> <ul style="list-style-type: none"> <li>• This command causes the router to become an ASBR by definition.</li> <li>• OSPF tags all routes learned through redistribution as external.</li> <li>• The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.</li> <li>• The metric is the cost you assign to the external route. The default is 20 for all protocols except BGP, whose default metric is 1.</li> <li>• The OSPF example redistributes BGP autonomous system 1, Level 1 routes into OSPF as Type 2 external routes.</li> <li>• The OSPFv3 example redistributes BGP autonomous system 1, Level 1 and 2 routes into OSPF. The external link type associated with the default route advertised into the OSPFv3 routing domain is the Type 1 external route.</li> </ul> <p><b>Note</b> RPL is not supported for OSPFv3.</p>

Command or Action	Purpose
<p><b>Step 5</b></p> <pre>summary-prefix address mask [not-advertise] [tag tag] or summary-prefix ipv6-prefix/prefix-length [not-advertise] [tag tag]</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0 or RP/0/RP0/CPU0:router(config-router)# summary-prefix 2010:11:22::/32</pre>	<p>(Optional) Creates aggregate addresses for OSPF.</p> <p>or</p> <p>(Optional) Creates aggregate addresses for OSPFv3.</p> <ul style="list-style-type: none"> <li>• This command provides external route summarization of the non-OSPF routes.</li> <li>• External ranges that are being summarized should be contiguous. Summarization of overlapping ranges from two different routers could cause packets to be sent to the wrong destination.</li> <li>• This command is optional. If you do not specify it, each route is included in the link-state database and advertised in LSAs.</li> <li>• In the OSPFv2 example, the summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external LSA.</li> <li>• In the OSPFv3 example, the summary address 2010:11:22::/32 has addresses such as 2010:11:22:0:1000::1, 2010:11:22:0:2000:679:1, and so on. Only the address 2010:11:22::/32 is advertised in the external LSA.</li> </ul>
<p><b>Step 6</b></p> <pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf)# end or RP/0/RP0/CPU0:router(config-ospf)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>



# Configuring OSPF Shortest Path First Throttling

This task explains how to configure SPF scheduling in millisecond intervals and potentially delay SPF calculations during times of network instability. This task is optional.

## Prerequisites

See the “[OSPF Shortest Path First Throttling](#)” section on page RC-186 for information about OSPF SPF throttling.

## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*  
or  
**router ospfv3** *process-name*
3. **router-id** {*router-id*}
4. **timers throttle spf** *spf-start spf-hold spf-max-wait*
5. **area** *area-id*
6. **interface** *type instance*
7. **end**  
or  
**commit**
8. **show ospf** [*process-name*]  
or  
**show ospfv3** [*process-name*]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i> or <b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1 or RP/0/RP0/CPU0:router(config)# router ospfv3 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  or  Enables OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.

	Command or Action	Purpose
Step 3	<p><b>router-id</b> <i>{router-id}</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf)# router-id  192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 4	<p><b>timers throttle spf</b> <i>spf-start spf-hold spf-max-wait</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf)# timers  throttle spf 10 4800 90000</p>	<p>Sets SPF throttling timers.</p>
Step 5	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures a backbone area.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 6	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar)# interface  POS 0/1/0/3</p>	<p>Enters interface configuration mode and associates one or more interfaces to the area.</p>

	Command or Action	Purpose
Step 3	<p><b>router-id</b> {<i>router-id</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 4	<p><b>timers throttle spf</b> <i>spf-start spf-hold spf-max-wait</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# timers throttle spf 10 4800 90000</p>	<p>Sets SPF throttling timers.</p>
Step 5	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures a backbone area.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 6	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3</p>	<p>Enters interface configuration mode and associates one or more interfaces to the area.</p>

	Command or Action	Purpose
Step 3	<p><b>router-id</b> {<i>router-id</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 4	<p><b>timers throttle spf</b> <i>spf-start spf-hold spf-max-wait</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# timers throttle spf 10 4800 90000</p>	<p>Sets SPF throttling timers.</p>
Step 5	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures a backbone area.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 6	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3</p>	<p>Enters interface configuration mode and associates one or more interfaces to the area.</p>

	Command or Action	Purpose
Step 7	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-ar-if)# end  or  RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 8	<pre>show ospf [process-name] or show ospfv3 [process-name]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show ospf 1  or  RP/0/RP0/CPU0:router# show ospfv3 2 </p>	<p>(Optional) Displays SPF throttling timers.</p>

## Examples

In the following example, the **show ospf EXEC** command is used to verify that the initial SPF schedule delay time, minimum hold time, and maximum wait time are configured correctly. Additional details are displayed about the OSPF process, such as the router type and redistribution of routes.

```
RP/0/RP0/CPU0:router# show ospf 1

Routing Process "ospf 1" with ID 192.168.4.3
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an autonomous system boundary router
  Redistributing External Routes from,
    ospf 2
  Initial SPF schedule delay 5 msec
  Minimum hold time between two consecutive SPF's 100 msec
  Maximum wait time between two consecutive SPF's 1000 msec
  Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
  Number of external LSA 0. Checksum Sum 00000000
  Number of opaque AS LSA 0. Checksum Sum 00000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  External flood list length 0
```

Non-Stop Forwarding enabled

**Note**

For a description of each output display field, see the **show ospf** command in the *OSPF Commands on Cisco IOS XR Software* module in the *Cisco IOS XR Routing Command Reference* document.

## Configuring Nonstop Forwarding Specific to Cisco for OSPF Version 2

This task explains how to configure OSPF NSF specific to Cisco on your NSF-capable router. This task is optional.

### Prerequisites

OSPF NSF requires that all neighbor networking devices be NSF aware, which happens automatically after you install the Cisco IOS XR image on the router. If an NSF-capable router discovers that it has non-NSF-aware neighbors on a particular network segment, it disables NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware routers continue to provide NSF capabilities.

See the “[Nonstop Forwarding for OSPF Version 2](#)” section on page RC-187 for conceptual information.

### Restrictions

The following are restrictions when configuring nonstop forwarding:

- OSPF Cisco NSF for virtual links is not supported.
- Neighbors must be NSF aware.

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **router-id** {*router-id*}
4. **nsf cisco**  
or  
**nsf cisco enforce global**
5. **nsf interval** *seconds*
6. **end**  
or  
**commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router ospf</b> <i>process-name</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1</p>	<p>Enables OSPF routing for the specified routing process and places the router in router configuration mode.</p> <p><b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.</p>
Step 3	<p><b>router-id</b> {<i>router-id</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 4	<p><b>nsf cisco</b> OR <b>nsf cisco enforce global</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# nsf cisco enforce global</p>	<p>Enables Cisco NSF operations for the OSPF process.</p> <ul style="list-style-type: none"> <li>Use the <b>nsf cisco</b> command without the optional <b>enforce</b> and <b>global</b> keywords to abort the NSF restart mechanism on the interfaces of detected non-NSF neighbors and allow NSF neighbors to function properly.</li> <li>Use the <b>nsf cisco</b> command with the optional <b>enforce</b> and <b>global</b> keywords if the router is expected to perform NSF during restart. However, if non-NSF neighbors are detected, NSF restart is canceled for the entire OSPF process.</li> </ul>

	Command or Action	Purpose
Step 5	<b>nsf interval</b> <i>seconds</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# nsf interval 120	Sets the minimum time between NSF restart attempts.  <b>Note</b> When you use this command, the OSPF process must be up for at least 90 seconds before OSPF attempts to perform an NSF restart.
Step 6	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# end or RP/0/RP0/CPU0:router(config-ospf)# commit	Saves configuration changes.  <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:   Uncommitted changes found, commit them before exiting (yes/no/cancel)?  [cancel]:   <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring OSPF Version 2 for MPLS Traffic Engineering

This task explains how to configure OSPF for MPLS TE. This task is optional.

For a description of the MPLS TE tasks and commands that allow you to configure the router to support tunnels, configure an MPLS tunnel that OSPF can use, and troubleshoot MPLS TE, see the *Implementing MPLS Traffic Engineering Configuration Guide*.

### Prerequisites

Your network must support the following Cisco IOS XR features before you enable MPLS TE for OSPF on your router:

- MPLS
- IP Cisco Express Forwarding (CEF)



#### Note

You must enter the commands in the following task on every OSPF router in the traffic-engineered portion of your network.



## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **router-id** {*router-id*}
4. **mpls traffic-eng router-id** {*ip-address* | *interface-type interface-instance*}
5. **area** *area-id*
6. **mpls traffic-eng**
7. **interface** *type instance*
8. **end**  
or  
**commit**
9. **show ospf** [*process-name*] [*area-id*] **mpls traffic-eng** {**link** | **fragment**}

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IPv4 address as the router ID.
Step 4	<b>mpls traffic-eng router-id</b> { <i>ip-address</i>   <i>interface-type interface-instance</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0	(Optional) Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.  <ul style="list-style-type: none"> <li>• This IP address is flooded to all nodes in TE LSAs.</li> <li>• For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router identifier of the destination node because that is the address that the traffic engineering topology database at the tunnel head uses for its path calculation.</li> <li>• We recommend that loopback interfaces be used for MPLS TE router ID because they are more stable than physical interfaces.</li> </ul>

	Command or Action	Purpose
Step 5	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0</p>	<p>Enters area configuration mode and configures an area for the OSPF process.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.</li> </ul>
Step 6	<p><b>mpls traffic-eng</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng</p>	<p>Configures the MPLS TE under the OSPF area.</p>
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface interface loopback0</p>	<p>Enters interface configuration mode and associates one or more interfaces to the area.</p>
Step 8	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 9	<p><b>show ospf</b> [<i>process-name</i>] [<i>area-id</i>] <b>mpls traffic-eng</b> {<b>link</b>   <b>fragment</b>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router# show ospf 1 0 mpls traffic-eng link</p>	<p>(Optional) Displays information about the links and fragments available on the local router for MPLS TE.</p>

## Examples

This section provides the following output examples:

- [Sample Output for the show ospf Command Before Configuring MPLS TE, page RC-229](#)
- [Sample Output for the show ospf mpls traffic-eng Command, page RC-229](#)
- [Sample Output for the show ospf Command After Configuring MPLS TE, page RC-230](#)

### Sample Output for the show ospf Command Before Configuring MPLS TE

In the following example, the **show route ospf EXEC** command verifies that POS interface 0/3/0/0 exists and MPLS TE is not configured:

```
RP/0/RP0/CPU0:router# show route ospf 1

O   11.0.0.0/24 [110/15] via 0.0.0.0, 3d19h, tunnel-te1
O   192.168.0.12/32 [110/11] via 11.1.0.2, 3d19h, POS0/3/0/0
O   192.168.0.13/32 [110/6] via 0.0.0.0, 3d19h, tunnel-te1
```

```
RP/0/RP0/CPU0:router#
```

### Sample Output for the show ospf mpls traffic-eng Command

In the following example, the **show ospf mpls traffic-eng EXEC** command verifies that the MPLS TE fragments are configured correctly:

```
RP/0/RP0/CPU0:router# show ospf 1 mpls traffic-eng fragment

OSPF Router with ID (192.168.4.3) (Process ID 1)

Area 0 has 1 MPLS TE fragment. Area instance is 3.
MPLS router address is 192.168.4.2
Next fragment ID is 1

Fragment 0 has 1 link. Fragment instance is 3.
Fragment has 0 link the same as last update.
Fragment advertise MPLS router address
Link is associated with fragment 0. Link instance is 3
  Link connected to Point-to-Point network
  Link ID :55.55.55.55
  Interface Address :192.168.50.21
  Neighbor Address :192.168.4.1
  Admin Metric :0
  Maximum bandwidth :19440000
  Maximum global pool reservable bandwidth :25000000
  Maximum sub pool reservable bandwidth   :3125000
  Number of Priority :8
  Global pool unreserved BW
  Priority 0 : 25000000 Priority 1 : 25000000
  Priority 2 : 25000000 Priority 3 : 25000000
  Priority 4 : 25000000 Priority 5 : 25000000
  Priority 6 : 25000000 Priority 7 : 25000000
  Sub pool unreserved BW
  Priority 0 : 3125000 Priority 1 : 3125000
  Priority 2 : 3125000 Priority 3 : 3125000
  Priority 4 : 3125000 Priority 5 : 3125000
  Priority 6 : 3125000 Priority 7 : 3125000
  Affinity Bit :0
```

In the following example, the **show ospf mpls traffic-eng EXEC** command verifies that the MPLS TE links on area instance 3 are configured correctly:

```
RP/0/RP0/CPU0:router# show ospf mpls traffic-eng link

      OSPF Router with ID (192.168.4.1) (Process ID 1)

Area 0 has 1  MPLS TE links. Area instance is 3.

Links in hash bucket 53.
Link is associated with fragment 0. Link instance is 3
Link connected to Point-to-Point network
Link ID :192.168.50.20
Interface Address :192.168.20.50
Neighbor Address :192.168.4.1
Admin Metric :0
Maximum bandwidth :19440000
Maximum global pool reservable bandwidth :25000000
Maximum sub pool reservable bandwidth   :3125000
Number of Priority :8
Global pool unreserved BW
Priority 0 : 25000000 Priority 1 : 25000000
Priority 2 : 25000000 Priority 3 : 25000000
Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
Priority 4 : 3125000 Priority 5 : 3125000
Priority 6 : 3125000 Priority 7 : 3125000
Affinity Bit :0
```

#### Sample Output for the show ospf Command After Configuring MPLS TE

In the following example, the **show route ospf EXEC** command verifies that the MPLS TE tunnels replaced POS interface 0/3/0/0 and that configuration was performed correctly:

```
RP/0/RP0/CPU0:router# show route ospf 1

O E2 192.168.10.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O E2 192.168.11.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O E2 192.168.1244.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O 192.168.12.0/24 [110/2] via 0.0.0.0, 00:00:15, tunnel2
```

## Configuring OSPFv3 Graceful Restart

This task explains how to configure a graceful restart for an OSPFv3 process. This task is optional.

### SUMMARY STEPS

1. **configure**
2. **router ospfv3 *process-name***
3. **graceful-restart**
4. **graceful-restart lifetime**
5. **graceful-restart interval <seconds>**
6. **graceful-restart helper**

7. **end**  
or  
**commit**
8. **show ospfv3** [*process-name* [*area-id*]] **database grace**

## DETAILED STEP

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospfv3</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospfv3 test	Enters router configuration mode for OSPFv3. The process name is a WORD that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters without spaces.
Step 3	<b>graceful-restart</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospfv3)#graceful-re start	Enables graceful restart on the current router.
Step 4	<b>graceful-restart lifetime</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart lifetime 120	Specifies a maximum duration for a graceful restart. <ul style="list-style-type: none"> <li>• The default lifetime is 95 seconds.</li> <li>• The range is 90 to 3600 seconds.</li> </ul>
Step 5	<b>graceful-restart interval</b> <i>&lt;seconds&gt;</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart interval 120	Specifies the interval (minimal time) between graceful restarts on the current router. <ul style="list-style-type: none"> <li>• The default value for the interval is 90 seconds.</li> <li>• The range is 90 to 3600 seconds.</li> </ul>
Step 6	<b>graceful-restart helper</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart helper disable	Disables the helper capability.

	Command or Action	Purpose
Step 7	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospfv3)# end  OR  RP/0/RP0/CPU0:router(config-ospfv3)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 8	<pre>show ospfv3 [process-name [area-id]] database grace</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show ospfv3 1 database grace</p>	<p>Displays the state of the graceful restart link.</p>

## Displaying Information About Graceful Restart

This section describes the tasks you can use to display information about a graceful restart.

- To see if the feature is enabled and when the last graceful restart ran, use the **show ospf** command. To see details for an OSPFv3 instance, use the **show ospfv3 process-name [area-id] database grace** command.

### Displaying the State of the Graceful Restart Feature

The following screen output shows the state of the graceful restart capability on the local router:

```
RP/0/RP0/CPU0:router# show ospfv3 1 database grace
```

```
Routing Process "ospfv3 1" with ID 2.2.2.2
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPFs 10000 msec
Maximum wait time between two consecutive SPFs 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 255
Number of external LSA 0. Checksum Sum 00000000
```

```

Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Graceful Restart enabled, last GR 11:12:26 ago (took 6 secs)
Area BACKBONE(0)
  Number of interfaces in this area is 1
  SPF algorithm executed 1 times
  Number of LSA 6. Checksum Sum 0x0268a7
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0

```

```
RP/0/RP0/CPU0:router#
```

### Displaying Graceful Restart Information for an OSPFv3 Instance

The following screen output shows the link state for an OSPFv3 instance:

```
RP/0/RP0/CPU0:router# show ospfv3 1 database grace
```

```

OSPFv3 Router with ID (2.2.2.2) (Process ID 1)

Router Link States (Area 0)
ADV Router      Age      Seq#      Fragment ID  Link count  Bits
1.1.1.1         1949    0x8000000e  0            1           1      None
2.2.2.2         2007    0x80000011  0            1           1      None

Link (Type-8) Link States (Area 0)
ADV Router      Age      Seq#      Link ID      Interface
1.1.1.1         180     0x80000006  1            PO0/2/0/0
2.2.2.2         2007    0x80000006  1            PO0/2/0/0

Intra Area Prefix Link States (Area 0)
ADV Router      Age      Seq#      Link ID      Ref-lstype  Ref-LSID
1.1.1.1         180     0x80000006  0            0x2001      0
2.2.2.2         2007    0x80000006  0            0x2001      0

Grace (Type-11) Link States (Area 0)
ADV Router      Age      Seq#      Link ID      Interface
2.2.2.2         2007    0x80000005  1            PO0/2/0/0

RP/0/RP0/CPU0:router#

```

## Configuring an OSPFv2 Sham Link

This task explains how to configure a provider edge (PE) router to establish an OSPFv2 sham link connection across a VPN backbone. This task is optional.

### Prerequisites

Before configuring a sham link in a Multiprotocol Label Switching (MPLS) VPN between provider edge (PE) routers, OSPF must be enabled as follows:

- Create an OSPF routing process.
- Configure a loopback interface that belongs to VRF and assign a IPv4 address with the host mask to it.
- Configure the sham link under the area submode.

See [“Enabling OSPF” section on page RC-194](#) for information on these OSPF configuration prerequisites.

## Summary Steps

1. **configure**
2. **interface** *type instance*
3. **vrf** *vrf-name*
4. **ipv4 address** *ipv4-address mask*
5. **end**
6. **router ospf** *instance-id*
7. **vrf** *vrf-name*
8. **router-id** {*router-id*}
9. **redistribute** **bgp** *process-id*
10. **area** *area-id*
11. **sham-link** *source-address destination-address*
12. **cost** *cost*
13. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# interface loopback 3	Enters interface configuration mode.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# vrf vrf1	Assigns an interface to the VPN routing and forwarding (VRF) instance.
Step 4	<b>ipv4 address</b> <i>ip-address mask</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38 255.255.255.225	Assigns an IP address and subnet mask to the interface.



	Command or Action	Purpose
Step 5	<p><b>end</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-if)# end</p>	<p>Saves configuration changes.</p> <p>When you issue the <b>end</b> command, the system prompts you to commit changes:</p> <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>• Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>• Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>• Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul>
Step 6	<p><b>router ospf</b> <i>instance-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf isp</p>	<p>Enables OSPF routing for the specified routing process, and places the router in router configuration mode. In this example, the OSPF instance is called isp.</p>
Step 7	<p><b>vrf</b> <i>vrf-name</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1</p>	<p>Creates a VRF instance and enters VRF configuration mode.</p>
Step 8	<p><b>router-id</b> {<i>router-id</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# router-id 192.168.4.3</p>	<p>Configures a router ID for the OSPF process.</p> <p><b>Note</b> We recommend using a stable IPv4 address as the router ID.</p>
Step 9	<p><b>redistribute bgp</b> <i>process-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# redistribute bgp 1</p>	<p>Redistributes OSPF routes from the one routing domain to another routing domain.</p> <ul style="list-style-type: none"> <li>• This command causes the router to become an ASBR by definition.</li> <li>• OSPF tags all routes learned through redistribution as external.</li> <li>• The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.</li> <li>• The BGP MED value is copied to the LSA metric field when BGP VPN routes are redistributed to OSPF.</li> </ul>
Step 10	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# area 0</p>	<p>Enters area configuration mode and configures an area for the OSPF process.</p> <ul style="list-style-type: none"> <li>• The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.</li> </ul>

	Command or Action	Purpose
Step 11	<p><b>sham-link</b> <i>source-address destination-address</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf-ar)#  sham-link 10.0.0.1 10.0.0.3</p>	Configures a point-to-point unnumbered interface between two VPN sites.
Step 12	<p><b>cost</b> <i>cost</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf-ar-sl)#  cost 76</p>	Explicitly specifies the cost of sending a packet on an OSPF interface. The specified cost overrides the auto-costing calculated default value for interfaces.
Step 13	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf-ar)# end  or  RP/0/RP0/CPU0:router(config-ospf-vrf-ar)#  commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Enabling Nonstop Routing for OSPFv2

This optional task describes how to enable nonstop routing (NSR) for OSPFv2 processes. NSR is disabled by default. When NSR is enabled, OSPF processes on active RP, synchronize all necessary data and states with the OSPF process on the standby RP. When the switchover happens, OSPF processes on newly active RP has all the necessary data and states to continue running and does not require any help from its neighbors.

### Summary Steps

- configure**
- router ospf** *instance-id*
- nsr**
- end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf isp	Enables OSPF routing for the specified routing process, and places the router in router configuration mode. In this example, the OSPF instance is called isp.
Step 3	<b>nsr</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# nsr	Enables NSR for the OSPFv2 process.
Step 4	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# end OR RP/0/RP0/CPU0:router(config-ospf)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Enabling Multicast-intact for OSPFv2

This optional task describes how to enable multicast-intact for OSPFv2 routes that use IPv4 addresses.

## Summary Steps

- configure**
- router ospf instance-id**
- mpls traffic-eng multicast-intact**

4. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf isp	Enables OSPF routing for the specified routing process, and places the router in router configuration mode. In this example, the OSPF instance is called isp.
Step 3	<b>mpls traffic-eng multicast-intact</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng multicast-intact	Enables multicast-intact.
Step 4	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# end or RP/0/RP0/CPU0:router(config-ospf)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Associating Interfaces to a VRF

This task explains how to associate an interface with a VPN Routing and Forwarding (VRF) instance.

### SUMMARY STEPS

1. **configure**

2. **router ospf** *process-name*
3. **vrf** *vrf-name*
4. **interface** *type instance*
5. **ipv4 address** *ip-address mask*
6. **ipv6 address** *ipv6-prefix/prefix-length [eui-64]*
7. **ipv4 mtu** *mtu*
8. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1	Creates a VRF instance and enters VRF configuration mode.
Step 4	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# interface POS 0/0/0/0	Enters interface configuration mode and associates one or more interfaces to the VRF.
Step 5	<b>ipv4 address</b> <i>ip-address mask</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38 255.255.255.224	Assigns an IP address and subnet mask to the interface.
Step 6	<b>ipv6 address</b> <i>ipv6-prefix/prefix-length [eui-64]</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# ipv6 address 2001:0DB8:C18:1::64	Specifies the IPv6 address assigned to the interface and enables IPV6 processing on the interface.  • A slash-mark (/) must precede the <i>prefix-length</i> argument, and there is no space between the <i>ipv6-prefix</i> argument and the slash.

	Command or Action	Purpose
Step 7	<pre>ipv4 mtu mtu</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-if)# ipv4 mtu 300</p>	Sets the maximum transmission unit (MTU) size of IPv4 packets sent on the interface.
Step 8	<pre>end or commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring OSPF as a Provider Edge to Customer Edge (PE-CE) Protocol

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **vrf** *vrf-name*
4. **router-id** {*router-id*}
5. **redistribute** *protocol* [*process-id*] {**level-1** | **level-1-2** | **level-2**} [**metric** *metric-value*] [**metric-type** *type-value*] [**match** {**internal** | **external** [{**1** | **2**} | **nssa-external** {**1** | **2**}] [**tag** *tag-value*] [**route-map** *map-tag* | **route-policy** *policy-tag*]
6. **area** *area-id*
7. **interface** *type instance*
8. **exit**
9. **domain-id** [**secondary**] **type** {**0005** | **0105** | **0205** | **8005**} **value** *value*
10. **domain-tag** *tag*
11. **disable-dn-bit-check**

```

12. end
    or
    commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1	Creates a VRF instance and enters VRF configuration mode.
Step 4	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# router-id 192.168.4.3	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IPv4 address as the router ID.
Step 5	<b>redistribute</b> <i>protocol</i> [ <i>process-id</i> ] { <b>level-1</b>   <b>level-1-2</b>   <b>level-2</b> } [ <b>metric</b> <i>metric-value</i> ] [ <b>metric-type</b> <i>type-value</i> ] [ <b>match</b> { <b>internal</b>   <b>external</b> [{ <b>1</b>   <b>2</b> }   <b>nssa-external</b> { <b>1</b>   <b>2</b> }] [ <b>tag</b> <i>tag-value</i> ] [ <b>route-map</b> <i>map-tag</i>   <b>route-policy</b> <i>policy-tag</i> ]}  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# redistribute bgp 1 level-1	Redistributes OSPF routes from one routing domain to another routing domain. <ul style="list-style-type: none"> <li>This command causes the router to become an ASBR by definition.</li> <li>OSPF tags all routes learned through redistribution as external.</li> <li>The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.</li> <li>The metric is the cost you assign to the external route. The default is 20 for all protocols except BGP, whose default metric is 1.</li> <li>The example shows the redistribution of BGP autonomous system 1, Level 1 routes into OSPF as Type 2 external routes.</li> </ul>
Step 6	<b>area</b> <i>area-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# area 0	Enters area configuration mode and configures an area for the OSPF process. <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.</li> </ul>

	Command or Action	Purpose
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf)#  interface POS 0/0/0/0</p>	Enters interface configuration mode and associates one or more interfaces to the VRF.
Step 8	<p><b>exit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-if)# exit</p>	Exits interface configuration mode.
Step 9	<p><b>domain-id</b> [<b>secondary</b>] <b>type</b> {0005   0105   0205   8005} <b>value</b> <i>value</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf)#  domain-id 0105 1AF234</p>	<p>Specifies the OSPF VRF domain ID.</p> <ul style="list-style-type: none"> <li>The <i>value</i> argument is a six-octet hex number.</li> </ul>
Step 10	<p><b>domain-tag</b> <i>tag</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf)#  domain-tag 234</p>	<p>Specifies the OSPF VRF domain tag.</p> <ul style="list-style-type: none"> <li>The valid range for <i>tag</i> is 0 to 4294967295.</li> </ul>
Step 11	<p><b>disable-dn-bit-check</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf)#  disable-dn-bit-check</p>	Specifies that down bits should be ignored.
Step 12	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-ospf-vrf)# end  or  RP/0/RP0/CPU0:router(config-ospf-vrf)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>



## Creating Multiple OSPF Instances (OSPF Process and a VRF)

This task explains how to create multiple OSPF instances. In this case, the instances are a normal OSPF instance and a VRF instance.

### SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **area** *area-id*
4. **interface** *type instance*
5. **exit**
6. **vrf** *vrf-name*
7. **area** *area-id*
8. **interface** *type instance*
9. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>area</b> <i>area-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode and configures a backbone area.  <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 4	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3	Enters interface configuration mode and associates one or more interfaces to the area.

	Command or Action	Purpose
Step 5	<code>exit</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# exit	Enters OSPF configuration mode.
Step 6	<code>vrf vrf-name</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1	Creates a VRF instance and enters VRF configuration mode.
Step 7	<code>area area-id</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# area 0	Enters area configuration mode and configures an area for a VRF instance under the OSPF process. <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.</li> </ul>
Step 8	<code>interface type instance</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-vrf)# interface POS 0/0/0/0	Enters interface configuration mode and associates one or more interfaces to the VRF.
Step 9	<code>end</code> OR <code>commit</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Multi-area Adjacency

This task explains how to create multiple areas on an OSPF primary interface.

## Restrictions

Multiple areas are created only on native point-to-point interfaces, such as Packet-over-SONET (PoS) or serial.

## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **area** *area-id*
4. **interface** *type instance*
5. **area** *area-id*
6. **multi-area-interface** *type instance*
7. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>area</b> <i>area-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode and configures a backbone area.  <ul style="list-style-type: none"> <li>• The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 4	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface POS 0/1/0/3	Enters interface configuration mode and associates one or more interfaces to the area.

	Command or Action	Purpose
Step 5	<p><b>area</b> <i>area-id</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1</p>	<p>Enters area configuration mode and configures an area used for multiple area adjacency.</p> <ul style="list-style-type: none"> <li>The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.</li> </ul>
Step 6	<p><b>multi-area-interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# multi-area-interface POS 0/1/0/3</p>	<p>Enables multiple adjacencies for different OSPF areas and enters multi-area interface configuration mode</p>
Step 7	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if-mif)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if-mif)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Label Distribution Protocol IGP Auto-configuration for OSPF

This task explains how to configure LDP auto-configuration for an OSPF instance.

Optionally, you can configure this feature for an area of an OSPF instance.

### SUMMARY STEPS

- configure**
- router ospf** *process-name*
- mpls ldp auto-config**
- end**  
OR  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf process-name</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>mpls ldp auto-config</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# mpls ldp auto-config	Enables LDP IGP interface auto-configuration for an OSPF instance.  <ul style="list-style-type: none"> <li>Optionally, this command can be configured for an area of an OSPF instance.</li> </ul>
Step 4	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# end OR RP/0/RP0/CPU0:router(config-ospf)# commit	Saves configuration changes.  <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:   <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Authentication Message Digest Management for OSPF

This task explains how to manage authentication of a keychain on the OSPF interface.

### Prerequisites

A valid keychain must be configured before this task can be attempted.

To learn how to configure a keychain and its associated attributes, see the “Implementing Key Chain Management on Cisco IOS XR Software” module of *Cisco IOS XR System Security Configuration Guide*.

## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **router-id** {*router-id*}
4. **area** *area-id*
5. **interface** *type instance*
6. **authentication message-digest keychain** *keychain*
7. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# router id 192.168.4.3	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IPv4 address as the router ID.
Step 4	<b>area</b> <i>area-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1	Enters area configuration mode.  The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.
Step 5	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet0/4/0/1	Enters interface configuration mode and associates one or more interfaces to the area.

Command or Action	Purpose
<p><b>Step 6</b></p> <pre>authentication message-digest keychain keychain</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication message-digest keychain ospf_int1</pre>	<p>Configures an MD5 keychain.</p> <p><b>Note</b> In the example, the <i>ospf_int1</i> keychain must be configured before you attempt this step.</p>
<p><b>Step 7</b></p> <pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar-if)# end or RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Examples

The following example shows how to configure the keychain *ospf\_intf\_1* that contains five key IDs. Each key ID is configured with different **send-lifetime** values; however, all key IDs specify the same text string for the key.

```
key chain ospf_intf_1
key 1
send-lifetime 11:30:30 May 1 2007 duration 600
cryptographic-algorithm MD5T
key-string clear ospf_intf_1
key 2
send-lifetime 11:40:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 3
send-lifetime 11:50:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 4
send-lifetime 12:00:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 5
send-lifetime 12:10:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
```

The following example shows that keychain authentication is enabled on the Gigabit Ethernet 0/4/0/1 interface:

```
RP/0/RP0/CPU0:router# show ospf 1 interface GigabitEthernet0/4/0/1

GigabitEthernet0/4/0/1 is up, line protocol is up
Internet Address 100.10.10.2/24, Area 0
Process ID 1, Router ID 2.2.2.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 2.2.2.1, Interface address 100.10.10.2
Backup Designated router (ID) 1.1.1.1, Interface address 100.10.10.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:02
Index 3/3, flood queue length 0
Next 0(0)/0(0)
Last flood scan length is 2, maximum is 16
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
Suppress hello for 0 neighbor(s)
Keychain-based authentication enabled
  Key id used is 3
Multi-area interface Count is 0
```

The following example shows output for configured keys that are active:

```
RP/0/RP0/CPU0:router# show key chain ospf_intf_1

Key-chain: ospf_intf_1/ -

Key 1 -- text "0700325C4836100B0314345D"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:30:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 2 -- text "10411A0903281B051802157A"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:40:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 3 -- text "06091C314A71001711112D5A"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:50:30, 01 May 2007 - (Duration) 600 [Valid now]
  Accept lifetime: Not configured
Key 4 -- text "151D181C0215222A3C350A73"
  cryptographic-algorithm -- MD5
  Send lifetime: 12:00:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 5 -- text "151D181C0215222A3C350A73"
  cryptographic-algorithm -- MD5
  Send lifetime: 12:10:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured

RP/0/RP0/CPU0:router#
```



# Configuring Generalized TTL Security Mechanism (GTSM) for OSPF

This task explains how to set the security time-to-live mechanism on an interface for GTSM.

## SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **router-id** {*router-id*}
4. **log adjacency changes** [detail] [disable]
5. **nsf** {cisco [enforce global] | ietf [helper disable]}
6. **timers throttle** *spf spf-start spf-hold spf-max-wait*
7. **area** *area-id*
8. **interface** *type instance*
9. **security ttl** [disable | hops *hop-count*]
10. **end**  
or  
**commit**
11. **show ospf** [*process-name*] [**vrf** *vrf-name*] [*area-id*] **interface** [*type instance*]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router ospf</b> <i>process-name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.  <b>Note</b> The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<b>router-id</b> { <i>router-id</i> }	Configures a router ID for the OSPF process.  <b>Note</b> We recommend using a stable IPv4 address as the router ID.
Step 4	<b>log adjacency changes</b> [detail] [enable   disable]  <b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# log adjacency changes detail	(Optional) Requests notification of neighbor changes.  <ul style="list-style-type: none"> <li>• By default, this feature is enabled.</li> <li>• The messages generated by neighbor changes are considered notifications, which are categorized as severity Level 5 in the <b>logging console</b> command. The <b>logging console</b> command controls which severity level of messages are sent to the console. By default, all severity level messages are sent.</li> </ul>

	Command or Action	Purpose
Step 5	<pre>nsf {cisco [enforce global]   ietf [helper disable]}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# nsf ietf</p>	(Optional) Configures NSF OSPF protocol. The example enables graceful restart.
Step 6	<pre>timers throttle spf spf-start spf-hold spf-max-wait</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# timers throttle spf 500 500 10000</p>	(Optional) Sets SPF throttling timers.
Step 7	<pre>area area-id</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf)# area 1</p>	Enters area configuration mode. The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.
Step 8	<pre>interface type instance</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet0/5/0/0</p>	Enters interface configuration mode and associates one or more interfaces to the area.
Step 9	<pre>security ttl [disable   hops hop-count]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-ospf-ar-if)# security ttl hops 2</p>	Sets the security TTL value in the IP header for OSPF packets.

	Command or Action	Purpose
Step 10	<pre>end or commit</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar-if)# end OR RP/0/RP0/CPU0:router(config-ospf-ar-if)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 11	<pre>show ospf [process-name] [vrf vrf-name] [area-id] interface [type instance]</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router# show ospf 1 interface GigabitEthernet0/5/0/0</pre>	<p>Displays OSPF interface information.</p>

## Examples

The following is sample output that displays the GTSM security TTL value configured on an OSPF interface:

```
RP/0/RP0/CPU0:router# show ospf 1 interface GigabitEthernet0/5/0/0

GigabitEthernet0/5/0/0 is up, line protocol is up
 Internet Address 120.10.10.1/24, Area 0
  Process ID 1, Router ID 100.100.100.100, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  TTL security enabled, hop count 2
  Designated Router (ID) 102.102.102.102, Interface address 120.10.10.3
  Backup Designated router (ID) 100.100.100.100, Interface address 120.10.10.1
  Flush timer for old DR LSA due in 00:02:36
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Index 1/1, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 102.102.102.102 (Designated Router)
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 0
RP/0/RP0/CPU0:router#
```

## Verifying OSPF Configuration and Operation

This task explains how to verify the configuration and operation of OSPF.

### SUMMARY STEPS

1. **show {ospf | ospfv3} [process-name]**
2. **show {ospf | ospfv3} [process-name] border-routers [router-id]**
3. **show {ospf | ospfv3}[process-name] database**
4. **show {ospf | ospfv3} [process-name] [area-id] flood-list interface type instance**
5. **show {ospf | ospfv3} [process-name] [vrf vrf-name] [area-id] interface [type instance]**
6. **show {ospf | ospfv3} [process-name] [area-id] neighbor [interface-type interface-instance] [neighbor-id] [detail]**
7. **clear {ospf | ospfv3}[process-name] process**
8. **clear {ospf | ospfv3}[process-name] statistics [neighbor [interface-type interface-instance] [ip-address]]**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>show {ospf   ospfv3} [process-name]</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show ospf group1	(Optional) Displays general information about OSPF routing processes.
Step 2	<b>show {ospf   ospfv3} [process-name] border-routers [router-id]</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show ospf group1 border-routers	(Optional) Displays the internal OSPF routing table entries to an ABR and ASBR.
Step 3	<b>show {ospf   ospfv3} [process-name] database</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show ospf group2 database	(Optional) Displays the lists of information related to the OSPF database for a specific router. <ul style="list-style-type: none"> <li>• The various forms of this command deliver information about different OSPF LSAs.</li> </ul>
Step 4	<b>show {ospf   ospfv3} [process-name] [area-id] flood-list interface type instance</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show ospf 100 flood-list interface pos 0/3/0/0	(Optional) Displays a list of OSPF LSAs waiting to be flooded over an interface.
Step 5	<b>show {ospf   ospfv3} [process-name] [vrf vrf-name] [area-id] interface [type instance]</b>  <b>Example:</b> RP/0/RP0/CPU0:router# show ospf 100 interface pos 0/3/0/0	(Optional) Displays OSPF interface information.

	Command or Action	Purpose
Step 6	<pre>show {ospf   ospfv3}[process-name] [area-id] neighbor [interface-type interface-instance] [neighbor-id] [detail]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show ospf 100 neighbor </p>	(Optional) Displays OSPF neighbor information on an individual interface basis.
Step 7	<pre>clear {ospf   ospfv3}[process-name] process</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear ospf 100 process </p>	(Optional) Resets an OSPF router process without stopping and restarting it.
Step 8	<pre>clear {ospf   ospfv3}[process-name] statistics [neighbor [interface-type interface-instance] [ip-address]]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# clear ospf 100 statistics </p>	(Optional) Clears the OSPF statistics of neighbor state transitions.

## Configuration Examples for Implementing OSPF on Cisco IOS XR Software

This section provides the following configuration examples:

- [Cisco IOS XR for OSPF Version 2 Configuration: Example, page RC-255](#)
- [CLI Inheritance and Precedence for OSPF Version 2: Example, page RC-257](#)
- [MPLS TE for OSPF Version 2: Example, page RC-258](#)
- [ABR with Summarization for OSPFv3: Example, page RC-258](#)
- [ABR Stub Area for OSPFv3: Example, page RC-258](#)
- [ABR Totally Stub Area for OSPFv3: Example, page RC-258](#)
- [Route Redistribution for OSPFv3: Example, page RC-259](#)
- [Virtual Link Configured Through Area 1 for OSPFv3: Example, page RC-259](#)
- [Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example, page RC-260](#)
- [VPN Backbone and Sham Link Configured for OSPF Version 2: Example, page RC-260](#)

### Cisco IOS XR for OSPF Version 2 Configuration: Example

The following example shows how an OSPF interface is configured for an area in Cisco IOS XR software.

In Cisco IOS XR software, area 0 must be explicitly configured with the **area** command and all interfaces that are in the range from 10.1.2.0 to 10.1.2.255 are bound to area 0. Interfaces are configured with the **interface** command (while the router is in area configuration mode) and the **area** keyword is not included in the interface statement.

**Cisco IOS XR Software Configuration**

```

interface POS 0/3/0/0
 ip address 10.1.2.1 255.255.255.255
 negotiation auto
!
router ospf 1
router-id 10.2.3.4
 area 0
   interface POS 0/3/0/0
!
!

```

The following example shows how OSPF interface parameters are configured for an area in Cisco IOS XR software.

In Cisco IOS XR software, OSPF interface-specific parameters are configured in interface configuration mode and explicitly defined for area 0. In addition, the **ip ospf** keywords are no longer required.

**Cisco IOS XR Software Configuration**

```

interface POS 0/3/0/0
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
router ospf 1
router-id 10.2.3.4
 area 0
   interface POS 0/3/0/0
     cost 77
     mtu-ignore
     authentication message-digest
     message-digest-key 1 md5 0 test
!
!

```

The following example shows the hierarchical CLI structure of Cisco IOS XR software:

In Cisco IOS XR software, OSPF areas must be explicitly configured, and interfaces configured under the area configuration mode are explicitly bound to that area. In this example, interface 10.1.2.0/24 is bound to area 0 and interface 10.1.3.0/24 is bound to area 1.

**Cisco IOS XR Software Configuration**

```

interface POS 0/3/0/0
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
interface POS 0/3/0/1
 ip address 10.1.3.1 255.255.255.0
 negotiation auto
!
router ospf 1
router-id 10.2.3.4
 area 0
   interface POS 0/3/0/0
!
 area 1
   interface POS 0/3/0/1
!
!

```

## CLI Inheritance and Precedence for OSPF Version 2: Example

The following example configures the cost parameter at different hierarchical levels of the OSPF topology, and illustrates how the parameter is inherited and how only one setting takes precedence. According to the precedence rule, the most explicit configuration is used.

The cost parameter is set to 5 in router configuration mode for the OSPF process. Area 1 sets the cost to 15 and area 6 sets the cost to 30. All interfaces in area 0 inherit a cost of 5 from the OSPF process because the cost was not set in area 0 or its interfaces.

In area 1, every interface has a cost of 15 because the cost is set in area 1 and 15 overrides the value 5 that was set in router configuration mode.

Area 4 does not set the cost, but POS interface 01/0/2 sets the cost to 20. The remaining interfaces in area 4 have a cost of 5 that is inherited from the OSPF process.

Area 6 sets the cost to 30, which is inherited by POS interfaces 0/1/0/3 and 0/2/0/3. POS interface 0/3/0/3 uses the cost of 1, which is set in interface configuration mode.

```
router ospf 1
router-id 10.5.4.3
cost 5
area 0
interface POS 0/1/0/0
!
interface POS 0/2/0/0
!
interface POS 0/3/0/0
!
!
area 1
cost 15
interface POS 0/1/0/1
!
interface POS 0/2/0/1
!
interface POS 0/3/0/1
!
!
area 4
interface POS 0/1/0/2
cost 20
!
interface POS 0/2/0/2
!
interface POS 0/3/0/2
!
!
area 6
cost 30
interface POS 0/1/0/3
!
interface POS 0/2/0/3
!
interface POS 0/3/0/3
cost 1
!
!
```

## MPLS TE for OSPF Version 2: Example

The following example shows how to configure the OSPF portion of MPLS TE. However, you still need to build an MPLS TE topology and create an MPLS TE tunnel. See the *Cisco IOS XR MPLS Configuration Guide* for information.

In this example, loopback interface 0 is associated with area 0 and MPLS TE is configured within area 0:

```
interface Loopback 0
  address 10.10.10.10 255.255.255.0
!
interface POS 0/2/0/0
  address 10.1.2.2 255.255.255.0
!
router ospf 1
  router-id 10.10.10.10
  nsf
  auto-cost reference-bandwidth 10000
  mpls traffic-eng router-id Loopback 0
  area 0
    mpls traffic-eng
      interface POS 0/2/0/0
      interface Loopback 0
```

## ABR with Summarization for OSPFv3: Example

The following example shows the prefix range 2300::/16 summarized from area 1 into the backbone:

```
router ospfv3 1
  router-id 192.168.0.217
  area 0
    interface POS 0/2/0/1
  area 1
    range 2300::/16
  interface POS 0/2/0/0
```

## ABR Stub Area for OSPFv3: Example

The following example shows that area 1 is configured as a stub area:

```
router ospfv3 1
  router-id 10.0.0.217
  area 0
    interface POS 0/2/0/1
  area 1
    stub
  interface POS 0/2/0/0
```

## ABR Totally Stub Area for OSPFv3: Example

The following example shows that area 1 is configured as a totally stub area:

```
router ospfv3 1
  router-id 10.0.0.217
  area 0
```



```
interface POS 0/2/0/1
area 1
stub no-summary
interface POS 0/2/0/0
```

## Route Redistribution for OSPFv3: Example

The following example uses prefix lists to limit the routes redistributed from other protocols.

Only routes with 9898:1000 in the upper 32 bits and with prefix lengths from 32 to 64 are redistributed from BGP 42. Only routes *not* matching this pattern are redistributed from BGP 1956.

```
ipv6 prefix-list list1
seq 10 permit 9898:1000::/32 ge 32 le 64

ipv6 prefix-list list2
seq 10 deny 9898:1000::/32 ge 32 le 64
seq 20 permit ::/0 le 128

router ospfv3 1
router-id 10.0.0.217
redistribute bgp 42
redistribute bgp 1956
distribute-list prefix-list list1 out bgp 42
distribute-list prefix-list list2 out bgp 1956
area 1
interface POS 0/2/0/0
```

## Virtual Link Configured Through Area 1 for OSPFv3: Example

This example shows how to set up a virtual link to connect the backbone through area 1 for the OSPFv3 topology that consists of areas 0 and 1 and virtual links 10.0.0.217 and 10.0.0.212:

### ABR 1 Configuration

```
router ospfv3 1
router-id 10.0.0.217
area 0
interface POS 0/2/0/1
area 1
virtual-link 10.0.0.212
interface POS 0/2/0/0
```

### ABR 2 Configuration

```
router ospfv3 1
router-id 10.0.0.212
area 0
interface POS 0/3/0/1
area 1
virtual-link 10.0.0.217
interface POS 0/2/0/0
```

## Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example

The following examples show how to configure a virtual link to your backbone and apply MD5 authentication. You must perform the steps described on both ABRs at each end of the virtual link.

After you explicitly configure the ABRs, the configuration is inherited by all interfaces bound to that area—unless you override the values and configure them explicitly for the interface.

To understand virtual links, see the [“Virtual Link and Transit Area for OSPF”](#) section on page RC-183.

In this example, all interfaces on router ABR1 use MD5 authentication:

```
router ospf ABR1
  router-id 10.10.10.10
  authentication message-digest
  message-digest-key 100 md5 0 cisco
  area 0
    interface pos 0/2/0/1
    interface pos 0/3/0/0
  area 1
    interface pos 0/3/0/1
    virtual-link 10.10.5.5
  !
!
```

In this example, only area 1 interfaces on router ABR3 use MD5 authentication:

```
router ospf ABR2
  router-id 10.10.5.5
  area 0
  area 1
    authentication message-digest
    message-digest-key 100 md5 0 cisco
    interface pos 0/9/0/1
    virtual-link 10.10.10.10
  area 3
    interface Loopback 0
    interface pos 0/9/0/0
  !
!
```

## VPN Backbone and Sham Link Configured for OSPF Version 2: Example

The following examples show how to configure a provider edge (PE) router to establish a VPN backbone and sham link connection:

```
logging console debugging
vrf vrf_1
  address-family ipv4 unicast
  import route-target
    100:1
  !
  export route-target
    100:1
  !
!
!
interface Loopback0
  ipv4 address 2.2.2.1 255.255.255.255
  !
interface Loopback1
  vrf vrf_1
  ipv4 address 10.0.1.3 255.255.255.255
```

```

    !
    interface GigabitEthernet0/2/0/2
      vrf vrf_1
      ipv4 address 100.10.10.2 255.255.255.0
    !
    interface GigabitEthernet0/2/0/3
      ipv4 address 100.20.10.2 255.255.255.0
    !
  !
route-policy pass-all
  pass
end-policy
!
router ospf 1
  log adjacency changes
  router-id 2.2.2.2
  vrf vrf_1
    router-id 22.22.22.2
    domain-id type 0005 value 111122223333
    domain-tag 140
    nsf ietf
    redistribute bgp 10
  area 0
    sham-link 10.0.1.3 10.0.0.101
  !
  interface GigabitEthernet0/2/0/2
  !
  !
  !
!
router ospf 2
  router-id 2.22.2.22
  area 0
    interface Loopback0
    !
    interface GigabitEthernet0/2/0/3
    !
  !
!
router bgp 10
  bgp router-id 2.2.2.1
  bgp graceful-restart restart-time 300
  bgp graceful-restart
  address-family ipv4 unicast
    redistribute connected
  !
  address-family vpnv4 unicast
  !
  neighbor 2.2.2.2
  remote-as 10
  update-source Loopback0
  address-family ipv4 unicast
  !
  address-family vpnv4 unicast
  !
!
vrf vrf_1
  rd 100:1
  address-family ipv4 unicast
    redistribute connected route-policy pass-all
    redistribute ospf 1 match internal external
  !
!
!
```

```

mpls ldp
  router-id 2.2.2.1
  interface GigabitEthernet0/2/0/3
  !
!
end

```

## Where to Go Next

To configure route maps through the RPL for OSPF Version 2, see the *Implementing Routing Policy on Cisco IOS XR Software* document.

To build an MPLS TE topology, create tunnels, and configure forwarding over the tunnel for OSPF Version 2; see the *Cisco IOS XR MPLS Configuration Guide*.

## Additional References

The following sections provide references related to implementing OSPF on Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
OSPF and OSPFv3 commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference</i>
MPLS TE feature information	<i>Implementing MPLS Traffic Engineering on Cisco IOS XR Software</i> module in <i>Cisco IOS XR MPLS Configuration Guide</i>

## Standards

Standards	Title
draft-ietf-ospf-multi-area-adj-07.txt	OSPF Multi-Area Adjacency
draft-ietf-pce-disco-proto-ospf-08.txt	OSPF Protocol Extensions for Path Computation Element (PCE)
draft-ietf-mpls-igp-sync-00.txt	LDP IGP Synchronization
draft-ietf-ospf-ospfv3-graceful-restart-07.txt	OSPFv3 Graceful Restart

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
RFC 1587	<i>The OSPF NSSA Option</i>
RFC 1793	<i>Extending OSPF to Support Demand Circuits</i>
RFC 2328	<i>OSPF Version 2</i>
RFC 2370	<i>The OSPF Opaque LSA Option</i>
RFC 2740	<i>OSPF for IPv6</i>
RFC 3101	<i>The OSPF Not-So-Stubby Area (NSSA) Option</i>
RFC 3137	<i>OSPF Stub Router Advertisement</i>
RFC 3509	<i>Alternative Implementations of OSPF Area Border Routers</i>
RFC 3623	<i>Graceful OSPF Restart</i>
RFC 3630	<i>Traffic Engineering (TE) Extensions to OSPF Version 2</i>
RFC 3682	<i>The Generalized TTL Security Mechanism (GTSM)</i>
RFC 3906	<i>Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels</i>
RFC 4136	<i>OSPF Refresh and Flooding Reduction in Stable Topologies</i>
RFC 4206	<i>Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)</i>
RFC 4124	<i>Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering</i>
RFC 4576	<i>Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) ownbit Extension for L3VPN</i>
RFC 4577	<i>OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)</i>
RFC 4750	<i>OSPF Version 2 Management Information Base</i>
RFC 4811	<i>OSPF Out-of-Band Link State Database (LSDB) Resynchronization</i>
RFC 4812	<i>OSPF Restart Signaling</i>
RFC 4813	<i>OSPF Link-Local Signaling</i>
RFC 4970	<i>Extensions to OSPF for Advertising Optional Router Capabilities</i>

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>



# Implementing IS-IS on Cisco IOS XR Software

Integrated Intermediate System-to-Intermediate System (IS-IS), Internet Protocol Version 4 (IPv4), is a standards-based Interior Gateway Protocol (IGP).

Cisco IOS XR implements the IP routing capabilities described in International Organization for Standardization (ISO)/International Engineering Consortium (IEC) 10589 and RFC 1995, and adds the standard extensions for single topology and multitopology IS-IS for IP Version 6 (IPv6).

This module describes the new and revised tasks you need to implement IS-IS (IPv4 and IPv6) on your Cisco IOS XR network.



## Note

For more information about IS-IS on Cisco IOS XR software and complete descriptions of the IS-IS commands listed in this module, you can refer to the [“Related Documents”](#) section of this module. To locate documentation for other commands that might appear while of executing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing IS-IS on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router. The ability to configure a broadcast medium connecting two networking devices as a point-to-point link was added.
Release 3.3.0	LDP IGP synchronization support was added for IPv4 unicast address families. See the <a href="#">“MPLS Label Distribution Protocol IGP Synchronization”</a> section on page RC-275 and <a href="#">“Configuring MPLS LDP IS-IS Synchronization”</a> section on page RC-306 for more information.  The <b>ispf startup-delay</b> command was removed from the <a href="#">“Setting SPF Interval for a Single-Topology IPv4 and IPv6 Configuration”</a> section on page RC-301.
Release 3.4.0	Support was added for the following features: <ul style="list-style-type: none"><li>• MPLS TE forwarding adjacency</li><li>• MPLS TE interarea tunnels</li><li>• Multicast topology</li></ul>

Release 3.5.0	Support was added for the following features: <ul style="list-style-type: none"> <li>• IS-IS keychain support for hello and LSP authentication and hitless key rollover</li> <li>• IP fast reroute loop-free alternates computation</li> <li>• Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) auto-configuration</li> </ul>
Release 3.6.0	No modification.

## Contents

- [Prerequisites for Implementing IS-IS on Cisco IOS XR Software, page RC-266](#)
- [Restrictions for Implementing IS-IS on Cisco IOS XR Software, page RC-266](#)
- [Information About Implementing IS-IS on Cisco IOS XR Software, page RC-266](#)
- [How to Implement IS-IS on Cisco IOS XR Software, page RC-277](#)
- [Configuration Examples for Implementing IS-IS on Cisco IOS XR Software, page RC-314](#)
- [Where to Go Next, page RC-316](#)
- [Additional References, page RC-317](#)

## Prerequisites for Implementing IS-IS on Cisco IOS XR Software

You must be in a user group associated with a task group that includes the proper task IDs for IS-IS commands. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of *Cisco IOS XR System Security Configuration Guide*.

## Restrictions for Implementing IS-IS on Cisco IOS XR Software

When multiple instances of IS-IS are being run, an interface can be associated with only one instance (process). Instances may not share an interface.

## Information About Implementing IS-IS on Cisco IOS XR Software

To implement IS-IS you need to understand the following concepts:

- [IS-IS Functional Overview, page RC-267](#)
- [Key Features Supported in the Cisco IOS XR IS-IS Implementation, page RC-267](#)
- [IS-IS Configuration Grouping, page RC-268](#)
- [IS-IS Configuration Modes, page RC-268](#)
- [Multitopology Configuration, page RC-269](#)
- [IPv6 Routing and Configuring IPv6 Addressing, page RC-269](#)
- [Limit LSP Flooding, page RC-269](#)



- [Maximum LSP Lifetime and Refresh Interval, page RC-270](#)
- [Overload Bit Configuration During Multitopology Operation, page RC-270](#)
- [Single-Topology IPv6 Support, page RC-271](#)
- [Multitopology IPv6 Support, page RC-271](#)
- [IS-IS Authentication, page RC-271](#)
- [Nonstop Forwarding, page RC-272](#)
- [Multi-Instance IS-IS, page RC-272](#)
- [Multiprotocol Label Switching Traffic Engineering, page RC-273](#)
- [Overload Bit on Router, page RC-273](#)
- [Default Routes, page RC-273](#)
- [Attached Bit on an IS-IS Instance, page RC-273](#)
- [IS-IS Support for Route Tags, page RC-274](#)
- [Multicast-Intact Feature, page RC-274](#)
- [Multicast Topology Support Using IS-IS, page RC-275](#)
- [MPLS Label Distribution Protocol IGP Synchronization, page RC-275](#)
- [Label Distribution Protocol IGP Auto-configuration, page RC-276](#)
- [MPLS TE Forwarding Adjacency, page RC-276](#)
- [MPLS TE Interarea Tunnels, page RC-276](#)
- [IP Fast Reroute, page RC-276](#)

## IS-IS Functional Overview

Small IS-IS networks are typically built as a single area that includes all routers in the network. As the network grows larger, it may be reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

The IS-IS routing protocol supports the configuration of backbone Level 2 and Level 1 areas and the necessary support for moving routing information between the areas. Routers establish Level 1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (interarea routing).

For Cisco IOS XR software, each IS-IS instance can support either a single Level 1 or Level 2 area, or one of each. By default, all IS-IS instances automatically support Level 1 and Level 2 routing. You can change the level of routing to be performed by a particular routing instance using the **is-type** command.

## Key Features Supported in the Cisco IOS XR IS-IS Implementation

The Cisco IOS XR implementation of IS-IS conforms to the IS-IS Version 2 specifications detailed in RFC 1195 and the IPv6 IS-IS functionality based on the Internet Engineering Task Force (IETF) IS-IS Working Group draft-ietf-isis-ipv6.txt document.

The following list outlines key features supported in the Cisco IOS XR implementation:

- Single topology IPv6
- Multitopology
- Nonstop forwarding (NSF), both Cisco proprietary and IETF
- Three-way handshake
- Mesh groups
- Multiple IS-IS instances
- Configuration of a broadcast medium connecting two networking devices as a point-to-point link
- Fast-flooding with different threads handling flooding and shortest path first (SPF).

**Note**

For information on IS-IS support for Bidirectional Forwarding Detection (BFD), see *Cisco IOS XR Interface and Hardware Configuration Guide* and *Cisco IOS XR Interface and Hardware Command Reference*.

## IS-IS Configuration Grouping

Cisco IOS XR groups all of the IS-IS configuration in router IS-IS configuration mode, including the portion of the interface configurations associated with IS-IS. To display the IS-IS configuration in its entirety, use the **show running router isis** command. The command output displays the running configuration for all configured IS-IS instances, including the interface assignments and interface attributes.

## IS-IS Configuration Modes

The following sections show how to enter each of the configuration modes. From a mode, you can enter the ? command to display the commands available in that mode.

### Router Configuration Mode

The following example shows how to enter router configuration mode:

```
RP/0/RP0/CPU0:router# configuration
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)#
```

### Router Address Family Configuration Mode

The following example shows how to enter router address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)#
```

### Interface Configuration Mode

The following example shows how to enter interface configuration mode:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface POS0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)#
```

## Interface Address Family Configuration Mode

The following example shows how to enter interface address family configuration mode:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface POS0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)#
```

## IS-IS Interfaces

IS-IS interfaces can be configured as one of the following types:

- active—advertises connected prefixes and forms adjacencies. This is the default for interfaces.
- passive—advertises connected prefixes but does not form adjacencies. The **passive** command is used to configure interfaces as passive. Passive interfaces should be used sparingly for important prefixes such as loopback addresses that need to be injected into the IS-IS domain. If many connected prefixes need to be advertised then the redistribution of connected routes with the appropriate policy should be used instead.
- suppressed—does not advertise connected prefixes but forms adjacencies. The **suppress** command is used to configure interfaces as suppressed.
- shutdown—does not advertise connected prefixes and does not form adjacencies. The **shutdown** command is used to disable interfaces without removing the IS-IS configuration.

## Multitopology Configuration

Cisco IOS XR software supports multitopology for IPv6 IS-IS unless single topology is explicitly configured in IPv6 address-family configuration mode.



### Note

IS-IS supports IP routing and not Open Systems Interconnection (OSI) Connectionless Network Service (CLNS) routing.

## IPv6 Routing and Configuring IPv6 Addressing

By default, IPv6 routing is disabled in the Cisco IOS XR software. To enable IPv6 routing, you must assign IPv6 addresses to individual interfaces in the router using the **ipv6 enable** or **ipv6 address** command. See the *Network Stack IPv4 and IPv6 Commands on Cisco IOS XR Software* module of *Cisco IOS XR IP Addresses and Services Command Reference*.

## Limit LSP Flooding

Limiting link-state packets (LSP) may be desirable in certain “meshy” network topologies. An example of such a network might be a highly redundant one such as a fully meshed set of point-to-point links over a nonbroadcast multiaccess (NBMA) transport. In such networks, full LSP flooding can limit network scalability. One way to restrict the size of the flooding domain is to introduce hierarchy by using multiple Level 1 areas and a Level 2 area. However, two other techniques can be used instead of or with hierarchy: Block flooding on specific interfaces and configure mesh groups.

Both techniques operate by restricting the flooding of LSPs in some fashion. A direct consequence is that although scalability of the network is improved, the reliability of the network (in the face of failures) is reduced because a series of failures may prevent LSPs from being flooded throughout the network, even though links exist that would allow flooding if blocking or mesh groups had not restricted their use. In such a case, the link-state databases of different routers in the network may no longer be synchronized. Consequences such as persistent forwarding loops can ensue. For this reason, we recommend that blocking or mesh groups be used only if specifically required, and then only after careful network design.

## Flood Blocking on Specific Interfaces

With this technique, certain interfaces are blocked from being used for flooding LSPs, but the remaining interfaces operate normally for flooding. This technique is simple to understand and configure, but may be more difficult to maintain and more error prone than mesh groups in the long run. The flooding topology that IS-IS uses is fine-tuned rather than restricted. Restricting the topology too much (blocking too many interfaces) makes the network unreliable in the face of failures. Restricting the topology too little (blocking too few interfaces) may fail to achieve the desired scalability.

To improve the robustness of the network in the event that all nonblocked interfaces drop, use the **csnp-interval** command in interface configuration mode to force periodic complete sequence number PDUs (CSNPs) packets to be used on blocked point-to-point links. The use of periodic CSNPs enables the network to become synchronized.

## Mesh Group Configuration

Configuring mesh groups (a set of interfaces on a router) can help to limit flooding. All routers reachable over the interfaces in a particular mesh group are assumed to be densely connected with each router having at least one link to every other router. Many links can fail without isolating one or more routers from the network.

In normal flooding, a new LSP is received on an interface and is flooded out over all other interfaces on the router. With mesh groups, when a new LSP is received over an interface that is part of a mesh group, the new LSP is not flooded over the other interfaces that are part of that mesh group.

## Maximum LSP Lifetime and Refresh Interval

By default, the router sends a periodic LSP refresh every 15 minutes. LSPs remain in a database for 20 minutes by default. If they are not refreshed by that time, they are deleted. You can change the LSP refresh interval or maximum LSP lifetime. The LSP interval should be less than the LSP lifetime or else LSPs time out before they are refreshed. In the absence of a configured refresh interval, the software adjusts the LSP refresh interval, if necessary, to prevent the LSPs from timing out.

## Overload Bit Configuration During Multitopology Operation

Because the overload bit applies to forwarding for a single topology, it may be configured and cleared independently for IPv4 and IPv6 during multitopology operation. For this reason, the overload is set from the router address family configuration mode. If the IPv4 overload bit is set, all routers in the area do not use the router for IPv4 transit traffic. However, they can still use the router for IPv6 transit traffic.

## Single-Topology IPv6 Support

Single-topology IPv6 support on Cisco IOS XR software allows IS-IS for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols, and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is used, only narrow link metrics, also known as old-style type, length, and value (TLV) arguments, may be employed. During single-topology operation, one shortest path first (SPF) computation for each level is used to compute both IPv4 and IPv6 routes. Using a single SPF is possible because both IPv4 IS-IS and IPv6 IS-IS routing protocols share a common link topology.

## Multitopology IPv6 Support

Multitopology IPv6 support on Cisco IOS XR software for IS-IS assumes that multitopology support is required as soon as it detects interfaces configured for both IPv6 and IPv4 within the IS-IS stanza.

Because multitopology is the default behavior in the software, you must explicitly configure IPv6 to use the same topology as IPv4 to enable single-topology IPv6. Configure the **single-topology** command in IPv6 router address family configuration submode of the IS-IS router stanza.

## IS-IS Authentication

Authentication is available to limit the establishment of adjacencies by using the **hello-password** command, and to limit the exchange of LSPs by using the **lsp-password** command.

IS-IS supports plain-text authentication, which does not provide security against unauthorized users. Plain-text authentication allows you to configure a password to prevent unauthorized networking devices from forming adjacencies with the router. The password is exchanged as plain text and is potentially visible to an agent able to view the IS-IS packets.

When an HMAC-MD5 password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

IS-IS stores a configured password using simple encryption. However, the plain-text form of the password is used in LSPs, sequence number protocols (SNPs), and hello packets, which would be visible to a process that can view IS-IS packets. The passwords can be entered in plain text (clear) or encrypted form.

To set the domain password, configure the **lsp-password** command for Level 2; to set the area password, configure the **lsp-password** command for Level 1.

The keychain feature allows IS-IS to reference configured keychains. IS-IS key chains enable hello and LSP keychain authentication. Keychains can be configured at the router level (in the case of the **lsp-password** command) and at the interface level (in the case of the **hello-password** command) within IS-IS. These commands reference the global keychain configuration and instruct the IS-IS protocol to obtain security parameters from the global set of configured keychains.

IS-IS is able to use the keychain to implement hitless key rollover for authentication. The key rollover specification is time based, and in the event of clock skew between the peers, the rollover process is impacted. The configurable tolerance specification allows for the accept window to be extended (before and after) by that margin. This accept window facilitates a hitless key rollover for applications (for example, routing and management protocols).

See *Cisco IOS XR System Security Guide* for information on keychain management.

## Nonstop Forwarding

On Cisco IOS XR software, NSF minimizes the amount of time a network is unavailable to its users following a route processor (RP) failover. The main objective of NSF is to continue forwarding IP packets and perform a graceful restart following an RP failover.

When a router restarts, all routing peers of that device usually detect that the device went down and then came back up. This transition results in what is called a *routing flap*, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. NSF helps to suppress routing flaps in NSF-aware devices, thus reducing network instability.

NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following an RP failover. When the NSF feature is configured, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards while the standby RP assumes control from the failed active RP during a failover. The ability of line cards to remain up through a failover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to NSF operation.

When the Cisco IOS XR router running IS-IS routing performs an RP failover, the router must perform two tasks to resynchronize its link-state database with its IS-IS neighbors. First, it must relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the link-state database for the network.

The IS-IS NSF feature offers two options when configuring NSF:

- IETF NSF
- Cisco NSF

If neighbor routers on a network segment are NSF aware, meaning that neighbor routers are running a software version that supports the IETF Internet draft for router restartability, they assist an IETF NSF router that is restarting. With IETF NSF, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a failover.

In Cisco IOS XR software, Cisco NSF checkpoints (stores persistently) all the state necessary to recover from a restart without requiring any special cooperation from neighboring routers. The state is recovered from the neighboring routers, but only using the standard features of the IS-IS routing protocol. This capability makes Cisco NSF suitable for use in networks in which other routers have not used the IETF standard implementation of NSF.

**Note**

If you configure IETF NSF on the Cisco IOS XR router and a neighbor router does not support IETF NSF, the affected adjacencies flap, but nonstop forwarding is maintained to all neighbors that do support IETF NSF. A restart reverts to a cold start if no neighbors support IETF NSF.

## Multi-Instance IS-IS

You may configure as many IS-IS instances as system resources (memory and interfaces) allow. Each interface may be associated with only a single IS-IS instance, and MPLS may be enabled for only a single IS-IS instance. Cisco IOS XR software prevents the double-booking of an interface by two instances at configuration time—two instances of MPLS configuration causes an error.

Because the Routing Information Base (RIB) treats each of the IS-IS instances as equal routing clients, you must be careful when redistributing routes between IS-IS instances. The RIB does not know to prefer Level 1 routes over Level 2 routes. For this reason, if you are running Level 1 and Level 2 instances, you must enforce the preference by configuring different administrative distances for the two instances.

## Multiprotocol Label Switching Traffic Engineering

The MPLS TE feature enables an MPLS backbone to replicate and expand the traffic engineering capabilities of Layer 2 ATM and Frame Relay networks. MPLS is an integration of Layer 2 and Layer 3 technologies.

For IS-IS, MPLS TE automatically establishes and maintains MPLS TE label-switched paths across the backbone by using Resource Reservation Protocol (RSVP). The route that a label-switched path uses is determined by the label-switched paths resource requirements and network resources, such as bandwidth. Available resources are flooded by using special IS-IS TLV extensions in the IS-IS. The label-switched paths are explicit routes and are referred to as traffic engineering (TE) tunnels.

## Overload Bit on Router

The overload bit is a special bit of state information that is included in an LSP of the router. If the bit is set on the router, it notifies routers in the area that the router is not available for transit traffic. This capability is useful in four situations:

1. During a serious but nonfatal error, such as limited memory.
2. During the startup and restart of the process. The overload bit can be set until the routing protocol has converged. However, it is not employed during a normal NSF restart or failover because doing so causes a routing flap.
3. During a trial deployment of a new router. The overload bit can be set until deployment is verified, then cleared.
4. During the shutdown of a router. The overload bit can be set to remove the router from the topology before the router is removed from service.

## Default Routes

You can force a default route into an IS-IS routing domain. Whenever you specifically configure redistribution of routes into an IS-IS routing domain, the Cisco IOS XR software does not, by default, redistribute the default route into the IS-IS routing domain. The **default-information originate** command generates a *default route* into IS-IS, which can be controlled by a route policy. You can use the route policy to identify the level into which the default route is to be announced, and you can specify other filtering options configurable under a route policy. You can use a route policy to conditionally advertise the default route, depending on the existence of another route in the routing table of the router.

## Attached Bit on an IS-IS Instance

The attached bit is set in a router that is configured with the **is-type** command and **level-1-2** keyword. The attached bit indicates that the router is connected to other areas (typically through the backbone). This functionality means that the router can be used by Level 1 routers in the area as the default route to

the backbone. The attached bit is usually set automatically as the router discovers other areas while computing its Level 2 SPF route. The bit is automatically cleared when the router becomes detached from the backbone.

**Note**

If the connectivity for the Level 2 instance is lost, the attached bit in the Level 1 instance LSP would continue sending traffic to the Level 2 instance and cause the traffic to be dropped.

To simulate this behavior when using multiple processes to represent the **level-1-2** keyword functionality, you would manually configure the attached bit on the Level 1 process.

## IS-IS Support for Route Tags

The IS-IS Support for route tags feature provides the capability to associate and advertise a tag with an IS-IS route prefix. Additionally, the feature allows you to prioritize the order of installation of route prefixes in the RIB based on a tag of a route. Route tags may also be used in route policy to match route prefixes (for example, to select certain route prefixes for redistribution).

## Multicast-Intact Feature

The multicast-intact feature provides the ability to run multicast routing (PIM) when IGP shortcuts are configured and active on the router. Both OSPFv2 and IS-IS support the multicast-intact feature. MPLS TE and IP multicast coexistence is supported in Cisco IOS XR software by using the **mpls traffic-eng multicast-intact** IS-IS or OSPF router command.

You can enable multicast-intact in the IGP when multicast routing protocols (PIM) are configured and IGP shortcuts are configured on the router. IGP shortcuts are MPLS tunnels that are exposed to IGP. The IGPs route the IP traffic over these tunnels to destinations that are downstream from the egress router of the tunnel (from an SPF perspective). PIM cannot use IGP shortcuts for propagating PIM joins because reverse path forwarding (RPF) cannot work across a unidirectional tunnel.

When you enable multicast-intact on an IGP, the IGP publishes a parallel or alternate set of equal-cost next-hops for use by PIM. These next-hops are called mcast-intact next-hops. The mcast-intact next-hops have the following attributes:

- They are guaranteed not to contain any IGP shortcuts.
- They are not used for unicast routing but are used only by PIM to look up an IPv4 next-hop to a PIM source.
- They are not published to the FIB.
- When multicast-intact is enabled on an IGP, all IPv4 destinations that were learned through link-state advertisements are published with a set equal-cost mcast-intact next-hops to the RIB. This attribute applies even when the native next-hops have no IGP shortcuts.
- In IS-IS, the max-paths limit is applied by counting both the native and mcast-intact next-hops together. (In OSPFv2, the behavior is slightly different.)



## Multicast Topology Support Using IS-IS

Multicast topology support allows for the configuration of IS-IS multicast topologies for IPv4 or IPv6 routing. IS-IS maintains a separate topology for multicast and runs a separate Shortest Path First (SPF) over the multicast topology. IS-IS multicast inserts routes from the IS-IS multicast topology into the multicast-unicast Routing Information Base (muRIB) table in the RIB for the corresponding address family. Since PIM uses the muRIB, PIM uses routes from the multicast topology instead of routes from the unicast topology.

## MPLS Label Distribution Protocol IGP Synchronization

Multiprotocol Label Switching (MPLS) Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) Synchronization ensures that LDP has completed label exchange before the IGP path is used for switching. MPLS traffic loss can occur in the following two situations:

- When an IGP adjacency is established, the router begins forwarding packets using the new adjacency before LDP has exchanged labels with peers on that link.
- When an LDP session closes, the router continues to forward traffic using the link associated with the LDP peer rather than using an alternate path with an established LDP session.

This feature provides a mechanism to synchronize LDP and IS-IS to minimize MPLS packet loss. The synchronization is accomplished by changing the link metric for a neighbor IS-IS link-state packet (LSP), based on the state of the LDP session.

When an IS-IS adjacency is established on a link but the LDP session is lost or LDP has not yet completed exchanging labels, IS-IS advertises the maximum metric on that link. In this instance, LDP IS-IS synchronization is not yet achieved.

**Note**

In IS-IS, a link with a maximum wide metric (0xFFFFFFFF) is not considered for shortest path first (SPF). Therefore, the maximum wide metric of -1 (0xFFFFFFFF) is used with MPLS LDP IGP synchronization.

When LDP IS-IS synchronization is achieved, IS-IS advertises a regular (configured or default) metric on that link.

## MPLS LDP-IGP Synchronization Compatibility with LDP Graceful Restart

LDP graceful restart protects traffic when an LDP session is lost. If a graceful restart-enabled LDP session fails, MPLS LDP IS-IS synchronization is still achieved on the interface while it is protected by graceful restart. MPLS LDP IGP synchronization is eventually lost under the following circumstances:

- LDP fails to restart before the LDP graceful restart reconnect timer expires.
- The LDP session on the protected interface fails to recover before the LDP graceful restart recovery timer expires.

## MPLS LDP-IGP Synchronization Compatibility with IGP Nonstop Forwarding

IS-IS nonstop forwarding (NSF) protects traffic during IS-IS process restarts and route processor (RP) failovers. LDP IS-IS synchronization is supported with IS-IS NSF only if LDP graceful restart is also enabled over the interface. If IS-IS NSF is not enabled, the LDP synchronization state is not retained across restarts and failovers.

## Label Distribution Protocol IGP Auto-configuration

Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) auto-configuration simplifies the procedure to enable LDP on a set of interfaces used by an IGP instance. LDP IGP auto-configuration can be used on a large number of interfaces (for example, when LDP is used for transport in the core) and on multiple IGP instances simultaneously.

This feature supports the IPv4 address family for the default VPN routing and forwarding (VRF) instance.

LDP IGP auto-configuration can also be explicitly disabled on individual interfaces under LDP using the **igp auto-config disable** command. This allows LDP to receive all IGP interfaces except the ones explicitly disabled.

See *Cisco IOS XR Multiprotocol Label Switching Configuration Guide* for information on configuring LDP IGP auto-configuration.

## MPLS TE Forwarding Adjacency

MPLS TE forwarding adjacency allows a network administrator to handle a traffic engineering, label switch path (LSP) tunnel as a link in an Interior Gateway Protocol (IGP) network, based on the Shortest Path First (SPF) algorithm. A forwarding adjacency can be created between routers in the same IS-IS level. The routers can be located multiple hops from each other. As a result, a TE tunnel is advertised as a link in an IGP network, with the cost of the link associated with it. Routers outside of the TE domain see the TE tunnel and use it to compute the shortest path for routing traffic throughout the network.

MPLS TE forwarding adjacency is considered in IS-IS SPF only if a two-way connectivity check is achieved. This is possible if the forwarding adjacency is bidirectional or the head end and tail end routers of the MPLS TE tunnel are adjacent.

The MPLS TE forwarding adjacency feature is supported by IS-IS. For details on configuring MPLS TE forwarding adjacency, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## MPLS TE Interarea Tunnels

MPLS TE interarea tunnels allow you to establish MPLS TE tunnels that span multiple IGP areas (Open Shortest Path First [OSPF]) and levels (IS-IS), removing the restriction that required that both the tunnel headend and tailend routers be in the same area. The IGP can be either IS-IS or OSPF. See the [“Configuring MPLS Traffic Engineering for IS-IS” section on page RC-295](#) for information on configuring MPLS TE for IS-IS.

For details on configuring MPLS TE interarea tunnels, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## IP Fast Reroute

The IP Fast Reroute (IPFRR) loop-free alternate (LFA) computation provides protection against link failure. Locally computed repair paths are used to prevent packet loss caused by loops that occur during network reconvergence after a failure. See IETF draft-ietf-rtgwg-ipfrr-framework-06.txt and draft-ietf-rtgwg-lf-conv-frmwk-00.txt for detailed information on IPFRR LFA.

**Note**

---

IPFRR is supported on the Cisco CRS-1 router.

---

IPFRR LFA is different from Multiprotocol Label Switching (MPLS) as it is applicable to networks using conventional IP routing and forwarding. See *Multiprotocol Label Switching Configuration Guide* for information on configuring MPLS IPFRR.

## How to Implement IS-IS on Cisco IOS XR Software

This section contains the following procedures:

- [Enabling IS-IS and Configuring Level 1 or Level 2 Routing, page RC-277](#) (required)
- [Configuring Single Topology for IS-IS, page RC-280](#) (optional)
- [Configuring Multitopology for IS-IS, page RC-284](#) (optional)
- [Controlling LSP Flooding for IS-IS, page RC-285](#) (optional)
- [Configuring Nonstop Forwarding for IS-IS, page RC-289](#) (optional)
- [Configuring Authentication for IS-IS, page RC-291](#) (optional)
- [Configuring Keychains for IS-IS, page RC-293](#) (optional)
- [Configuring MPLS Traffic Engineering for IS-IS, page RC-295](#) (optional)
- [Tuning Adjacencies for IS-IS, page RC-298](#) (optional)
- [Setting SPF Interval for a Single-Topology IPv4 and IPv6 Configuration, page RC-301](#) (optional)
- [Customizing Routes for IS-IS, page RC-303](#) (optional)
- [Configuring MPLS LDP IS-IS Synchronization, page RC-306](#) (optional)
- [Enabling Multicast-Intact, page RC-308](#) (optional)
- [Tagging IS-IS Interface Routes, page RC-309](#) (optional)
- [Setting the Priority for Adding Prefixes to the RIB, page RC-311](#) (optional)
- [Configuring IP Fast Reroute Loop-free Alternate, page RC-313](#) (optional)

**Note**

---

To save configuration changes, you must commit changes when the system prompts you.

---

## Enabling IS-IS and Configuring Level 1 or Level 2 Routing

This task explains how to enable IS-IS and configure the routing level for an area.

**Note**

---

Configuring the routing level in Step 4 is optional, but is highly recommended to establish the proper level of adjacencies.

---

### Prerequisites

Although you can configure IS-IS before you configure an IP address, no IS-IS routing occurs until at least one IP address is configured.

## SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **net** *network-entity-title*
4. **is-type** { **level-1** | **level-1-2** | **level-2-only** }
5. **end**  
or  
**commit**
6. **show isis** [**instance** *instance-id*] **protocol**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>• By default, all IS-IS instances are automatically Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul>
Step 3	<b>net</b> <i>network-entity-title</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00	Configures network entity titles (NETs) for the routing instance. <ul style="list-style-type: none"> <li>• Specify a NET for each routing instance if you are configuring multi-instance IS-IS.</li> <li>• This example configures a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110.00.</li> <li>• To specify more than one area address, specify additional NETs. Although the area address portion of the NET differs, the systemID portion of the NET must match exactly for all of the configured items.</li> </ul>

	Command or Action	Purpose
Step 4	<p><b>is-type</b> {<b>level-1</b>   <b>level-1-2</b>   <b>level-2-only</b>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only</p>	<p>(Optional) Configures the system type (area or backbone router).</p> <ul style="list-style-type: none"> <li>• By default, every IS-IS instance acts as a <b>level-1-2</b> router.</li> <li>• The <b>level-1</b> keyword configures the software to perform Level 1 (intra-area) routing only. Only Level 1 adjacencies are established. The software learns about destinations inside its area only. Any packets containing destinations outside the area are sent to the nearest <b>level-1-2</b> router in the area.</li> <li>• The <b>level-2-only</b> keyword configures the software to perform Level 2 (backbone) routing only, and the router establishes only Level 2 adjacencies, either with other Level 2-only routers or with <b>level-1-2</b> routers.</li> <li>• The <b>level-1-2</b> keyword configures the software to perform both Level 1 and Level 2 routing. Both Level 1 and Level 2 adjacencies are established. The router acts as a border router between the Level 2 backbone and its Level 1 area.</li> </ul>
Step 5	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)# end OR  RP/0/RP0/CPU0:router(config-isis)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 6	<p><b>show isis</b> [<b>instance</b> <i>instance-id</i>] <b>protocol</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis protocol</p>	<p>(Optional) Displays summary information about the IS-IS instance.</p>

## Configuring Single Topology for IS-IS

After an IS-IS instance is enabled, it must be configured to compute routes for a specific network topology.

This task explains how to configure the operation of the IS-IS protocol on an interface for an IPv4 or IPv6 topology.

### Restrictions

To enable the router to run in single-topology mode, configure each of the IS-IS interfaces with all of the address families enabled and “single-topology” in the address-family IPv6 unicast in the IS-IS router stanza. You can use either the IPv6 address family or both IPv4 and IPv6 address families, but your configuration must represent the set of all active address families on the router. Additionally, explicitly enable single-topology operation by configuring it in the IPv6 router address family submode.

Two exceptions to these instructions exist:

1. If the address-family stanza in the IS-IS process contains the **adjacency-check disable** command, then an interface is not required to have the address family enabled.
2. The **single-topology** command is not valid in the ipv4 address-family submode.

The default metric style for single topology is narrow metrics. However, you can use either wide metrics or narrow metrics. How to configure them depends on how single topology is configured. If both IPv4 and IPv6 are enabled and single topology is configured, the metric style is configured in the **address-family ipv4** stanza. You may configure the metric style in the **address-family ipv6** stanza, but it is ignored in this case. If only IPv6 is enabled and single topology is configured, then the metric style is configured in the **address-family ipv6** stanza.

### SUMMARY STEPS

1. **configure**
2. **interface** *type number*
3. **ipv4 address** *address mask*  
or  
**ipv6 address** *ipv6-prefix/prefix-length [eui-64]*  
or  
**ipv6 address** *ipv6-address {/prefix-length | link-local}*  
or  
**ipv6 enable**
4. **exit**
5. **router isis** *instance-id*
6. **net** *network-entity-title*
7. **address-family ipv6 [unicast]**
8. **single-topology**
9. **exit**
10. **interface** *type instance*
11. **circuit-type** {**level-1** | **level-1-2** | **level-2-only**}
12. **address-family** {**ipv4** | **ipv6**} [**unicast** | **multicast**]

13. **end**  
or  
**commit**
14. **show isis** [**instance** *instance-id*] **interface** [*type instance*] [**detail**] [**level** {**1** | **2**}]
15. **show isis** [**instance** *instance-id*] **topology** [**systemid** *system-id*] [**level** {**1** | **2**}] [**summary**]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>interface</b> <i>type number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# interface POS 0/1/0/3	Enters interface configuration mode.

Command or Action	Purpose
<p><b>Step 3</b></p> <pre> <b>ipv4 address</b> <i>address mask</i> or <b>ipv6 address</b> <i>ipv6-prefix/prefix-length [eui-64]</i> or <b>ipv6 address</b> <i>ipv6-address {/prefix-length   link-local}</i> or <b>ipv6 enable</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.1.3 255.255.255.0 or RP/0/RP0/CPU0:router(config-if)# ipv6 address 3ffe:1234:c18:1::/64 eui-64 or RP/0/RP0/CPU0:router(config-if)# ipv6 address FE80::260:3EFF:FE11:6770 link-local or RP/0/RP0/CPU0:router(config-if)# ipv6 enable </pre>	<p>Defines the IPv4 address for the interface. An IP address is required on all interfaces in an area enabled for IS-IS if any one interface is configured for IS-IS routing.</p> <p>or</p> <p>Specifies an IPv6 network assigned to the interface and enables IPv6 processing on the interface with the <b>eui-64</b> keyword.</p> <p>or</p> <p>Specifies an IPv6 address assigned to the interface and enables IPv6 processing on the interface with the <b>link-local</b> keyword.</p> <p>or</p> <p>Automatically configures an IPv6 link-local address on the interface while also enabling the interface for IPv6 processing.</p> <ul style="list-style-type: none"> <li>• The link-local address can be used only to communicate with nodes on the same link.</li> <li>• Specifying the <b>ipv6 address</b> <i>ipv6-prefix/prefix-length</i> interface configuration command without the <b>eui-64</b> keyword configures site-local and global IPv6 addresses.</li> <li>• Specifying the <b>ipv6 address</b> <i>ipv6-prefix/prefix-length</i> command with the <b>eui-64</b> keyword configures site-local and global IPv6 addresses with an interface ID in the low-order 64 bits of the IPv6 address. Only the 64-bit network prefix for the address needs to be specified; the last 64 bits are automatically computed from the interface ID.</li> <li>• Specifying the <b>ipv6 address</b> command with the <b>link-local</b> keyword configures a link-local address on the interface that is used instead of the link-local address that is automatically configured when IPv6 is enabled on the interface.</li> </ul>
<p><b>Step 4</b></p> <pre> <b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-if)# exit </pre>	<p>Exits interface configuration mode, and returns the router to global configuration mode.</p>
<p><b>Step 5</b></p> <pre> <b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp </pre>	<p>Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.</p> <ul style="list-style-type: none"> <li>• By default, all IS-IS instances are Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>



	Command or Action	Purpose
Step 6	<p><b>net</b> <i>network-entity-title</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)# net  47.0004.004d.0001.0001.0c11.1110.00</p>	<p>Configures NETs for the routing instance.</p> <ul style="list-style-type: none"> <li>Specify a NET for each routing instance if you are configuring multi-instance IS-IS. You can specify a name for a NET and for an address.</li> <li>This example configures a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110.00.</li> <li>To specify more than one area address, specify additional NETs. Although the area address portion of the NET differs, the system ID portion of the NET must match exactly for all of the configured items.</li> </ul>
Step 7	<p><b>address-family ipv6</b> [<b>unicast</b>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)#  address-family ipv6 unicast</p>	<p>Specifies the IPv6 address family and enters router address family configuration mode.</p> <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 8	<p><b>single-topology</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-af)#  single-topology</p>	<p>(Optional) Configures the link topology for IPv4 when IPv6 is configured.</p> <ul style="list-style-type: none"> <li>The <b>single-topology</b> command is valid only in IPv6 submode. The command instructs IPv6 to use the single topology rather than the default configuration of a separate topology in the multitopology mode.</li> <li>See the “<a href="#">Single-Topology IPv6 Support</a>” section on <a href="#">page RC-271</a> for more information.</li> </ul>
Step 9	<p><b>exit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-af)# exit</p>	<p>Exits router address family configuration mode, and returns the router to router configuration mode.</p>
Step 10	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)# interface  POS 0/1/0/3</p>	<p>Enters interface configuration mode.</p>
Step 11	<p><b>circuit-type</b> {<b>level-1</b>   <b>level-1-2</b>   <b>level-2-only</b>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  circuit-type level-1-2</p>	<p>(Optional) Configures the type of adjacency.</p> <ul style="list-style-type: none"> <li>The default circuit type is the configured system type (configured through the <b>is-type</b> command).</li> <li>Typically, the circuit type must be configured when the router is configured as only <b>level-1-2</b> and you want to constrain an interface to form only <b>level-1</b> or <b>level-2-only</b> adjacencies.</li> </ul>
Step 12	<p><b>address-family</b> {<b>ipv4</b>   <b>ipv6</b>} [<b>unicast</b>   <b>multicast</b>]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  address-family ipv6 unicast</p>	<p>Specifies the IPv4 or IPv6 address family, and enters interface address family configuration mode.</p> <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family on the interface.</li> </ul>

	Command or Action	Purpose
Step 13	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if-af)# end  OR  RP/0/RP0/CPU0:router(config-isis-if-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 14	<pre>show isis [instance instance-id] interface [type instance] [detail] [level {1   2}]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis interface  POS0/1/0/1</p>	(Optional) Displays information about the IS-IS interface.
Step 15	<pre>show isis [instance instance-id] topology [systemid system-id] [level {1   2}] [summary]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis topology</p>	(Optional) Displays a list of connected routers in all areas.

## Configuring Multitopology for IS-IS

Multitopology is configured in the same way as the single topology for IPv4 and IPv6 address families. However, the **single-topology** command is omitted, invoking the default multitopology behavior. This task is optional.

## Controlling LSP Flooding for IS-IS

Flooding of LSPs can limit network scalability. You can control LSP flooding by tuning your LSP database parameters on the router globally or on the interface. This task is optional.

Many of the commands to control LSP flooding contain an option to specify the level to which they apply. Without the option, the command applies to both levels. If an option is configured for one level, the other level continues to use the default value. To configure options for both levels, use the command twice. For example:

```
RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 1200 level 2
RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 1100 level 1
```

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **lsp-refresh-interval** *seconds* [**level** {**1** | **2**}]
4. **lsp-check-interval** *seconds* [**level** {**1** | **2**}]
5. **lsp-gen-interval** {[**initial-wait** *initial* | **secondary-wait** *secondary* | **maximum-wait** *maximum*] ...}[**level** {**1** | **2**}]
6. **lsp-mtu** *bytes* [**level** {**1** | **2**}]
7. **max-lsp-lifetime** *seconds* [**level** {**1** | **2**}]
8. **ignore-lsp-errors** **disable**
9. **interface** *type instance*
10. **lsp-interval** *milliseconds* [**level** {**1** | **2**}]
11. **csnp-interval** *seconds* [**level** {**1** | **2**}]
12. **retransmit-interval** *seconds* [**level** {**1** | **2**}]
13. **retransmit-throttle-interval** *milliseconds* [**level** {**1** | **2**}]
14. **mesh-group** {*number* | **blocked**}
15. **end**  
or  
**commit**
16. **show isis interface** [*type instance* | **level** {**1** | **2**}] [**brief**]
17. **show isis** [**instance** *instance-id*] **database** [**level** {**1** | **2**}] [**detail** | **summary** | **verbose**] [**\*** | *lsp-id*]
18. **show isis** [**instance** *instance-id*] **lsp-log** [**level** {**1** | **2**}]
19. **show isis database-log** [**level** {**1** | **2**}]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul>
Step 3	<b>lsp-refresh-interval seconds [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 10800	(Optional) Sets the time between regeneration of LSPs that contain different sequence numbers <ul style="list-style-type: none"> <li>The refresh interval should always be set lower than the <b>max-lsp-lifetime</b> command.</li> </ul>
Step 4	<b>lsp-check-interval seconds [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-check-interval 240	(Optional) Configures the time between periodic checks of the entire database to validate the checksums of the LSPs in the database. <ul style="list-style-type: none"> <li>This operation is costly in terms of CPU and so should be configured to occur infrequently.</li> </ul>
Step 5	<b>lsp-gen-interval</b> {[ <b>initial-wait initial</b>   <b>secondary-wait secondary</b>   <b>maximum-wait maximum</b> ] ...} [ <b>level {1   2}</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5	(Optional) Reduces the rate of LSP generation during periods of instability in the network. Helps reduce the CPU load on the router and number of LSP transmissions to its IS-IS neighbors. <ul style="list-style-type: none"> <li>During prolonged periods of network instability, repeated recalculation of LSPs can cause an increased CPU load on the local router. Further, the flooding of these recalculated LSPs to the other Intermediate Systems in the network causes increased traffic and can result in other routers having to spend more time running route calculations.</li> </ul>
Step 6	<b>lsp-mtu bytes [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-mtu 1300	(Optional) Sets the maximum transmission unit (MTU) size of LSPs.
Step 7	<b>max-lsp-lifetime seconds [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# max-lsp-lifetime 11000	(Optional) Sets the initial lifetime given to an LSP originated by the router. <ul style="list-style-type: none"> <li>This is the amount of time that the LSP persists in the database of a neighbor unless the LSP is regenerated or refreshed.</li> </ul>

	Command or Action	Purpose
Step 8	<p><b>ignore-lsp-errors</b> <b>disable</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)#  ignore-lsp-errors disable</p>	(Optional) Sets the router to purge LSPs received with checksum errors.
Step 9	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis)# interface  POS 0/1/0/3</p>	Enters interface configuration mode.
Step 10	<p><b>lsp-interval</b> <i>milliseconds</i> [<b>level</b> {<b>1</b>   <b>2</b>}]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  lsp-interval 100</p>	(Optional) Configures the amount of time between each LSP sent on an interface.
Step 11	<p><b>csnp-interval</b> <i>seconds</i> [<b>level</b> {<b>1</b>   <b>2</b>}]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  csnp-interval 30 level 1</p>	<p>(Optional) Configures the interval at which periodic CSNP packets are sent on broadcast interfaces.</p> <ul style="list-style-type: none"> <li>• Sending more frequent CSNPs means that adjacent routers must work harder to receive them.</li> <li>• Sending less frequent CSNP means that differences in the adjacent routers may persist longer.</li> </ul>
Step 12	<p><b>retransmit-interval</b> <i>seconds</i> [<b>level</b> {<b>1</b>   <b>2</b>}]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  retransmit-interval 60</p>	(Optional) Configures the amount of time that the sending router waits for an acknowledgment before it considers that the LSP was not received and subsequently resends.
Step 13	<p><b>retransmit-throttle-interval</b> <i>milliseconds</i> [<b>level</b> {<b>1</b>   <b>2</b>}]</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  retransmit-throttle-interval 1000</p>	<p>(Optional) Configures the amount of time between retransmissions on each LSP on a point-to-point interface.</p> <ul style="list-style-type: none"> <li>• This time is usually greater than or equal to the <b>lsp-interval</b> command time because the reason for lost LSPs may be that a neighboring router is busy. A longer interval gives the neighbor more time to receive transmissions.</li> </ul>
Step 14	<p><b>mesh-group</b> {<i>number</i>   <b>blocked</b>}</p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  mesh-group blocked</p>	<p>(Optional) Optimizes LSP flooding in NBMA networks with highly meshed, point-to-point topologies.</p> <ul style="list-style-type: none"> <li>• This command is appropriate only for an NBMA network with highly meshed, point-to-point topologies.</li> </ul>

	Command or Action	Purpose
Step 15	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)# end  OR  RP/0/RP0/CPU0:router(config-isis-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 16	<pre>show isis interface [type instance   level {1   2}] [brief]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis interface  POS0/1/0/1 brief</p>	(Optional) Displays information about the IS-IS interface.
Step 17	<pre>show isis [instance instance-id] database [level {1   2}] [detail   summary   verbose] [*   lsp-id]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis database  level 1</p>	(Optional) Displays the IS-IS LSP database.
Step 18	<pre>show isis [instance instance-id] lsp-log [level {1   2}]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis lsp-log</p>	(Optional) Displays LSP log information.
Step 19	<pre>show isis database-log [level {1   2}]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis database-log  level 1</p>	(Optional) Display IS-IS database log information.

## Configuring Nonstop Forwarding for IS-IS

This task explains how to configure your router with NSF that allows the Cisco IOS XR software to resynchronize the IS-IS link-state database with its IS-IS neighbors after a process restart. The process restart could be due to an:

- RP failover (for a warm restart)
- Simple process restart (due to an IS-IS reload or other administrative request to restart the process)
- IS-IS software upgrade

In all cases, NSF mitigates link flaps and loss of user sessions. This task is optional.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **nsf** { **cisco** | **ietf** }
4. **nsf interface-expires** *number*
5. **nsf interface-timer** *seconds*
6. **nsf lifetime** *seconds*
7. **end**  
or  
**commit**
8. **show running-config** [*command*]

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>• You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul>
Step 3	<b>nsf</b> { <b>cisco</b>   <b>ietf</b> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# nsf ietf	Enables NSF on the next restart. <ul style="list-style-type: none"> <li>• Enter the <b>cisco</b> keyword to run IS-IS in heterogeneous networks that might not have adjacent NSF-aware networking devices.</li> <li>• Enter the <b>ietf</b> keyword to enable IS-IS in homogeneous networks where <i>all</i> adjacent networking devices support IETF draft-based restartability.</li> </ul>

	Command or Action	Purpose
Step 4	<p><b>nsf interface-expires</b> <i>number</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# nsf interface-expires 1</p>	<p>Configures the number of resends of an acknowledged NSF-restart acknowledgment.</p> <ul style="list-style-type: none"> <li>If the resend limit is reached during the NSF restart, the restart falls back to a cold restart.</li> </ul>
Step 5	<p><b>nsf interface-timer</b> <i>seconds</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis) nsf interface-timer 15</p>	<p>Configures the number of seconds to wait for each restart acknowledgment.</p>
Step 6	<p><b>nsf lifetime</b> <i>seconds</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# nsf lifetime 20</p>	<p>Configures the maximum route lifetime following an NSF restart.</p> <ul style="list-style-type: none"> <li>This command should be configured to the length of time required to perform a full NSF restart because it is the amount of time that the Routing Information Base (RIB) retains the routes during the restart.</li> <li>Setting this value too high results in stale routes.</li> <li>Setting this value too low could result in routes purged too soon.</li> </ul>
Step 7	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# end OR RP/0/RP0/CPU0:router(config-isis)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 8	<p><b>show running-config</b> [<i>command</i>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router# show running-config router isis isp</p>	<p>(Optional) Displays the entire contents of the currently running configuration file or a subset of that file.</p> <ul style="list-style-type: none"> <li>Verify that “nsf” appears in the IS-IS configuration of the NSF-aware device.</li> <li>This example shows the contents of the configuration file for the “isp” instance only.</li> </ul>



## Configuring Authentication for IS-IS

This task explains how to configure authentication for IS-IS. This task is optional.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **lsp-password** {**hmac-md5** | **text**} {**clear** | **encrypted**} *password* [**level** {**1** | **2**}] [**send-only**] [**snp send-only**]
4. **interface** *type instance*
5. **hello-password** {**hmac-md5** | **text**} {**clear** | **encrypted**} *password* [**level** {**1** | **2**}] [**send-only**]
6. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# <b>configure</b>	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# <b>router isis</b> <i>isp</i>	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>• You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>

Command or Action	Purpose
<p><b>Step 3</b></p> <pre>lsp-password {hmac-md5   text} {clear   encrypted} password [level {1   2}] [send-only] [snp send-only]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1</p>	<p>Configures the LSP authentication password.</p> <ul style="list-style-type: none"> <li>The <b>hmac-md5</b> keyword specifies that the password is used in HMAC-MD5 authentication.</li> <li>The <b>text</b> keyword specifies that the password uses cleartext password authentication.</li> <li>The <b>clear</b> keyword specifies that the password is unencrypted when entered.</li> <li>The <b>encrypted</b> keyword specifies that the password is encrypted using a two-way algorithm when entered.</li> <li>The <b>level 1</b> keyword sets a password for authentication in the area (in Level 1 LSPs and Level SNPs).</li> <li>The <b>level 2</b> keywords set a password for authentication in the backbone (the Level 2 area).</li> <li>The <b>send-only</b> keyword adds authentication to LSP and sequence number protocol data units (SNPs) when they are sent. It does not authenticate received LSPs or SNPs.</li> <li>The <b>snp send-only</b> keyword adds authentication to SNPs when they are sent. It does not authenticate received SNPs.</li> </ul> <p><b>Note</b> To disable SNP password checking, the <b>snp send-only</b> keywords must be specified in the <b>lsp-password</b> command.</p>
<p><b>Step 4</b></p> <pre>interface type instance</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3</p>	<p>Enters interface configuration mode.</p>

	Command or Action	Purpose
Step 5	<pre>hello-password {hmac-md5   text} {clear   encrypted} password [level {1   2}] [send-only]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  hello-password text clear mypassword</p>	Configures the authentication password for an IS-IS interface.
Step 6	<pre>end</pre> <p>or</p> <pre>commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)# end  or  RP/0/RP0/CPU0:router(config-isis-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Keychains for IS-IS

This task explains how to configure keychains for IS-IS. This task is optional.

Keychains can be configured at the router level (**lsp-password** command) and at the interface level (**hello-password** command) within IS-IS. These commands reference the global keychain configuration and instruct the IS-IS protocol to obtain security parameters from the global set of configured keychains. The router-level configuration (**lsp-password** command) sets the keychain to be used for all IS-IS LSPs generated by this router, as well as for all Sequence Number Protocol Data Units (SN PDUs). The keychain used for HELLO PDUs is set at the interface level, and may be set differently for each interface configured for IS-IS.

### SUMMARY STEPS

- configure**
- router isis** *instance-id*
- lsp-password keychain** *keychain-name* [level {1 | 2}] [send-only] [snp send-only]
- interface** *type instance*
- hello-password keychain** *keychain-name* [level {1 | 2}] [send-only]

```

6. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>
Step 3	<b>lsp-password keychain</b> <i>keychain-name</i> [ <b>level</b> {1   2}] [ <b>send-only</b> ] [ <b>snp send-only</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# lsp-password keychain isis_a level 1	Configures the keychain.
Step 4	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3	Enters interface configuration mode.

	Command or Action	Purpose
Step 5	<pre>hello-password keychain keychain-name [level {1   2}] [send-only]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)#  hello-password keychain isis_b</p>	Configures the authentication password for an IS-IS interface.
Step 6	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if)# end  or  RP/0/RP0/CPU0:router(config-isis-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring MPLS Traffic Engineering for IS-IS

This task explains how to configure IS-IS for MPLS TE. This task is optional.

For a description of the MPLS TE tasks and commands that allow you to configure the router to support tunnels, configure an MPLS tunnel that IS-IS can use, and troubleshoot MPLS TE, see *Implementing MPLS Traffic Engineering on Cisco IOS XR Software*.

### Prerequisite

Your network must support the MPLS Cisco IOS XR software feature before you enable MPLS TE for IS-IS on your router.



#### Note

You must enter the commands in the following task list on every IS-IS router in the traffic-engineered portion of your network.

### Restrictions

MPLS traffic engineering currently does not support routing and signaling of LSPs over unnumbered IP links. Therefore, do not configure the feature over those links.

## SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **address-family** {*ipv4* | *ipv6*} [**unicast** | **multicast**]
4. **mpls traffic-eng level** {**1** | **2**}
5. **mpls traffic-eng router-id** {*ip-address* | *interface-name interface-instance*}
6. **metric-style wide** [level {**1** | **2**}]
7. **end**  
or  
**commit**
8. **show isis** [**instance** *instance-id*] **mpls traffic-eng tunnel**
9. **show isis** [**instance** *instance-id*] **mpls traffic-eng adjacency-log**
10. **show isis** [**instance** *instance-id*] **mpls traffic-eng advertisements**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul>
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } [ <b>unicast</b>   <b>multicast</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 4	<b>mpls traffic-eng level</b> { <b>1</b>   <b>2</b> }	Configures a router running IS-IS to flood MPLS TE link information into the indicated IS-IS level.
Step 5	<b>mpls traffic-eng router-id</b> { <i>ip-address</i>   <i>interface-name interface-instance</i> }  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# mpls traffic-eng router-id loopback0	Specifies that the MPLS TE router identifier for the node is the given IP address or an IP address associated with the given interface.

	Command or Action	Purpose
Step 6	<pre>metric-style wide [level {1   2}]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-af)#  metric-style wide level 1</p>	Configures a router to generate and accept only wide link metrics in the Level 1 area.
Step 7	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-af)# end</p> <p>OR</p> <pre>RP/0/RP0/CPU0:router(config-isis-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?</pre> <pre>[cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 8	<pre>show isis [instance instance-id] mpls traffic-eng tunnel</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis instance isp mpls traffic-eng tunnel</p>	(Optional) Displays MPLS TE tunnel information.
Step 9	<pre>show isis [instance instance-id] mpls traffic-eng adjacency-log</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis instance isp mpls traffic-eng adjacency-log</p>	(Optional) Displays a log of MPLS TE IS-IS adjacency changes.
Step 10	<pre>show isis [instance instance-id] mpls traffic-eng advertisements</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis instance isp mpls traffic-eng advertisements</p>	(Optional) Displays the latest flooded record from MPLS TE.

## Tuning Adjacencies for IS-IS

This task explains how to enable logging of adjacency state changes, alter the timers for IS-IS adjacency packets, and display various aspects of adjacency state. Tuning your IS-IS adjacencies increases network stability when links are congested. This task is optional.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, which means that the level modifiers are meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the specification of the level options.

The options configurable in the interface submode apply only to that interface. By default, the values are applied to both Level 1 and Level 2.

The **hello-password** command can be used to prevent adjacency formation with unauthorized or undesired routers. This ability is particularly useful on a LAN, where connections to routers with which you have no desire to establish adjacencies are commonly found.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **log adjacency changes**
4. **interface** *type number*
5. **hello-padding** { **disable** | **sometimes** } [**level** { **1** | **2** }]
6. **hello-interval** *seconds* [**level** { **1** | **2** }]
7. **hello-multiplier** *multiplier* [**level** { **1** | **2** }]
8. **hello-password** { **hmac-md5** | **text** } { **clear** | **encrypted** } *password* [**level** { **1** | **2** }] [**send-only**]
9. **end**  
or  
**commit**
10. **show isis** [**instance** *instance-id*] **adjacency** [*interface-type interface-instance*] [**detail**] [**systemid** *system-id*]
11. **show isis adjacency-log**
12. **show isis** [**instance** *instance-id*] **interface** [*type instance*] [**brief** | **detail**] [**level** { **1** | **2** }]
13. **show isis** [**instance** *instance-id*] **neighbors** [*interface-type interface-instance*] [**summary**] [**detail**] [**systemid** *system-id*]



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>
Step 3	<b>log adjacency changes</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# log adjacency changes	Generates a log message when an IS-IS adjacency changes state (up or down).
Step 4	<b>interface type number</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3	Enters interface configuration mode.
Step 5	<b>hello-padding {disable   sometimes} [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# hello-padding sometimes	Configures padding on IS-IS hello PDUs for an IS-IS interface on the router. <ul style="list-style-type: none"> <li>Hello padding applies to only this interface and not to all interfaces.</li> </ul>
Step 6	<b>hello-interval seconds [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 6	Specifies the length of time between hello packets that the software sends.
Step 7	<b>hello-multiplier multiplier [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# hello-multiplier 10	Specifies the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down. <ul style="list-style-type: none"> <li>A higher value increases the networks tolerance for dropped packets, but also may increase the amount of time required to detect the failure of an adjacent router.</li> <li>Conversely, not detecting the failure of an adjacent router can result in greater packet loss.</li> </ul>

	Command or Action	Purpose
Step 8	<pre>hello-password {hmac-md5   text} {clear   encrypted} password [level {1   2}] [send-only]</pre> <p><b>Example:</b> RP/0/RP0/CPU1:router(config-isis-if)# hello-password text clear mypassword</p>	Specifies that this system include authentication in the hello packets and requires successful authentication of the hello packet from the neighbor to establish an adjacency.
Step 9	<pre>end</pre> <p>or</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# end or RP/0/RP0/CPU0:router(config-isis-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 10	<pre>show isis [instance instance-id] adjacency [interface-type interface-instance] [detail] [systemid system-id]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show isis instance isp adjacency ipv4</p>	(Optional) Displays IS-IS adjacencies.
Step 11	<pre>show isis adjacency-log</pre> <p><b>Example:</b> RP/0/RP0/CPU1:router# show isis adjacency-log</p>	(Optional) Displays a log of the most recent adjacency state transitions.

	Command or Action	Purpose
Step 12	<pre>show isis [instance instance-id] interface [type instance] [brief   detail] [level {1   2}]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show isis interface POS 0/1/0/1 brief</p>	(Optional) Displays information about the IS-IS interface.
Step 13	<pre>show isis [instance instance-id] neighbors [interface-type interface-instance] [summary] [detail] [systemid system-id]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show isis neighbors summary</p>	(Optional) Displays information about IS-IS neighbors.

## Setting SPF Interval for a Single-Topology IPv4 and IPv6 Configuration

This task explains how to make adjustments to the SPF calculation to tune router performance. This task is optional.

Because the SPF calculation computes routes for a particular topology, the tuning attributes are located in the router address family configuration submode. SPF calculation computes routes for Level 1 and Level 2 separately.

When IPv4 and IPv6 address families are used in a single-topology mode, only a single SPF for the IPv4 topology exists. The IPv6 topology “borrows” the IPv4 topology; therefore, no SPF calculation is required for IPv6. To tune the SPF calculation parameters for single-topology mode, configure the **address-family ipv4 unicast** command.

The incremental SPF algorithm can be enabled separately. When enabled, the incremental shortest path first (ISPF) is not employed immediately. Instead, the full SPF algorithm is used to “seed” the state information required for the ISPF to run. The startup delay prevents the ISPF from running for a specified interval after an IS-IS restart (to permit the database to stabilize). After the startup delay elapses, the ISPF is principally responsible for performing all of the SPF calculations. The reseed interval enables a periodic running of the full SPF to ensure that the iSFP state remains synchronized.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **address-family** {*ipv4* | *ipv6*} [**unicast** | **multicast**]
4. **spf-interval** {[**initial-wait** *initial* | **secondary-wait** *secondary* | **maximum-wait** *maximum*]  
...}[**level** {**1** | **2**}]
5. **ispf** [**level** {**1** | **2**}]
6. **end**  
or  
**commit**
7. **show isis** [**instance** *instance-id*] [[**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** | **safi-all**]] **spf-log** [**level** {**1** | **2**}] [**ispf** | **fspf** | **prc** | **nhc**] [**detail** | **verbose**] [**last** *number* | **first** *number*]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul>
Step 3	<b>address-family {ipv4   ipv6} [unicast   multicast]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 4	<b>spf-interval</b> {[initial-wait <i>initial</i>   secondary-wait <i>secondary</i>   maximum-wait <i>maximum</i> ] ...} [level {1   2}]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 30	(Optional) Controls the minimum time between successive SPF calculations. <ul style="list-style-type: none"> <li>This value imposes a delay in the SPF computation after an event trigger and enforces a minimum elapsed time between SPF runs.</li> <li>If this value is configured too low, the router can lose too many CPU resources when the network is unstable.</li> <li>Configuring the value too high delays changes in the network topology that result in lost packets.</li> <li>The SPF interval does not apply to the running of the ISPF because that algorithm runs immediately on receiving a changed LSP.</li> </ul>
Step 5	<b>ispf [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# ispf	(Optional) Configures incremental IS-IS ISPF to calculate network topology.

	Command or Action	Purpose
Step 6	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-af)# end  OR  RP/0/RP0/CPU0:router(config-isis-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 7	<pre>show isis [instance instance-id] [[ipv4   ipv6   afi-all] [unicast   multicast   safi-all]] spf-log [level {1   2}] [ispf   fspf   prc   nhc] [detail   verbose] [last number   first number]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# show isis instance 1 spf-log ipv6</p>	<p>(Optional) Displays how often and why the router has run a full SPF calculation.</p>

## Customizing Routes for IS-IS

This task explains how to perform route functions that include injecting default routes into your IS-IS routing domain and redistributing routes learned in another IS-IS instance. This task is optional.

### SUMMARY STEPS

- configure**
- router isis** *instance-id*
- set-overload-bit** [**on-startup** {*delay* | **wait-for-bgp**}] [**level** {**1** | **2**}]
- address-family** {**ipv4** | **ipv6**} [**unicast** | **multicast**]
- default-information originate** [**route-policy** *route-policy-name*]
- redistribute isis** *instance* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric*] [**metric-type** {**internal** | **external**}] [**policy** *policy-name*]
- summary-prefix** *address/prefix-length* [**level** {**1** | **2**}]  
or  
**summary-prefix** *ipv6-prefix/prefix-length* [**level** {**1** | **2**}]

8. **maximum-paths** *route-number*
9. **distance** *weight* [*address/prefix-length* [*route-list-name*]]
10. **set-attached-bit**
11. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. <ul style="list-style-type: none"> <li>By default, all IS-IS instances are automatically Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>
Step 3	<b>set-overload-bit</b> [ <b>on-startup</b> { <i>delay</i>   <i>wait-for-bgp</i> }] [ <b>level</b> { <b>1</b>   <b>2</b> }]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# set-overload-bit	(Optional) Sets the overload bit. <b>Note</b> The configured overload bit behavior does not apply to NSF restarts because the NSF restart does not set the overload bit during restart.
Step 4	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } [ <b>unicast</b>   <b>multicast</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 5	<b>default-information originate</b> [ <b>route-policy</b> <i>route-policy-name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# default-information originate	(Optional) Injects a default IPv4 or IPv6 route into an IS-IS routing domain. <ul style="list-style-type: none"> <li>The <b>route-policy</b> keyword and <i>route-policy-name</i> argument specify the conditions under which the IPv4 or IPv6 default route is advertised.</li> <li>If the <b>route-policy</b> keyword is omitted, then the IPv4 or IPv6 default route is unconditionally advertised at Level 2.</li> </ul>
Step 6	<b>redistribute isis</b> <i>instance</i> [ <b>level-1</b>   <b>level-2</b>   <b>level-1-2</b> ] [ <b>metric</b> <i>metric</i> ] [ <b>metric-type</b> { <b>internal</b>   <b>external</b> }] [ <b>policy</b> <i>policy-name</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# redistribute isis 2 level-1	(Optional) Redistributes routes from one IS-IS instance into another instance. <ul style="list-style-type: none"> <li>In this example, an IS-IS instance redistributes Level 1 routes from another IS-IS instance.</li> </ul>

Command or Action	Purpose
<p><b>Step 7</b></p> <pre>summary-prefix address/prefix-length [level {1   2}] or summary-prefix ipv6-prefix/prefix-length [level {1   2}]</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 10.1.0.0/16 level 1 or RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 3003:xxx::/24 level 1</pre>	<p>(Optional) Allows a Level 1-2 router to summarize Level 1 IPv4 and IPv6 prefixes at Level 2, instead of advertising the Level 1 prefixes directly when the router advertises the summary.</p> <ul style="list-style-type: none"> <li>This example specifies an IPv4 address and mask.</li> </ul> <p>or</p> <ul style="list-style-type: none"> <li>This example specifies an IPv6 prefix, and the command must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.</li> <li>Note that IPv6 prefixes must be configured only in the IPv6 router address family configuration submode, and IPv4 prefixes in the IPv4 router address family configuration submode.</li> </ul>
<p><b>Step 8</b></p> <pre>maximum-paths route-number</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-isis-af)# maximum-paths 16</pre>	<p>(Optional) Configures the maximum number of parallel paths allowed in a routing table.</p>
<p><b>Step 9</b></p> <pre>distance weight [address/prefix-length [route-list-name]]</pre> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-isis-af)# distance 90</pre>	<p>(Optional) Defines the administrative distance assigned to routes discovered by the IS-IS protocol.</p> <ul style="list-style-type: none"> <li>A different administrative distance may be applied for IPv4 and IPv6.</li> </ul>

	Command or Action	Purpose
Step 10	<b>set-attached-bit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# set-attached-bit	(Optional) Configures an IS-IS instance with an attached bit in the Level 1 LSP.
Step 11	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# end or RP/0/RP0/CPU0:router(config-isis-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring MPLS LDP IS-IS Synchronization

This task explains how to enable Multiprotocol Label Switching (MPLS) Label Distribution Protocol (LDP) IS-IS synchronization. MPLS LDP synchronization can be enabled for an address family under interface configuration mode. Only IPv4 unicast address family is supported. This task is optional.

### SUMMARY STEPS

- configure**
- router isis** *instance-id*
- interface** *type instance*
- address-family ipv4 unicast**
- mpls ldp sync** [level {1 | 2}]
- end**  
or  
**commit**



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router isis instance-id</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp</p>	<p>Enables IS-IS routing for the specified routing process, and places the router in router configuration mode.</p> <ul style="list-style-type: none"> <li>By default, all IS-IS instances are automatically Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> command.</li> </ul>
Step 3	<p><b>interface type number</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3</p>	Enters interface configuration mode.
Step 4	<p><b>address-family ipv4 unicast</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast</p>	Specifies the IPv4 address family and enters router address family configuration mode.
Step 5	<p><b>mpls ldp sync [level {1   2}]</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if-af)# mpls ldp sync level 1</p>	Enables MPLS LDP synchronization for the IPv4 address family under interface POS 0/1/0/3.
Step 6	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if-af)# end OR RP/0/RP0/CPU0:router(config-isis-if-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Enabling Multicast-Intact

This optional task describes how to enable multicast-intact for IS-IS routes that use IPv4 and IPv6 addresses.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **address-family** {*ipv4* | *ipv6*} [**unicast** | **multicast**]
4. **mpls traffic-eng multicast-intact**
5. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. In this example, the IS-IS instance is called isp.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } [ <b>unicast</b>   <b>multicast</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>• This example specifies the unicast IPv6 address family.</li> </ul>

	Command or Action	Purpose
Step 4	<pre>mpls traffic-eng multicast-intact</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# mpls traffic-eng multicast-intact</p>	Enables multicast-intact.
Step 5	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# end OR RP/0/RP0/CPU0:router(config-isis-af)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Tagging IS-IS Interface Routes

This optional task describes how to associate a tag with a connected route of an IS-IS interface.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **address-family** {ipv4 | ipv6} [unicast | multicast]
4. **metric-style wide** [transition] [level {1 | 2}]
5. **exit**
6. **interface** *type number*
7. **address-family** {ipv4 | ipv6} [unicast | multicast]
8. **tag** *tag*
9. **end**  
or  
**commit**
10. **show isis** [ipv4 | ipv6 | afi-all] [unicast | multicast | safi-all] route [detail]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. In this example, the IS-IS instance is called isp.
Step 3	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } [ <b>unicast</b>   <b>multicast</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 4	<b>metric-style wide</b> [ <b>transition</b> ] [ <b>level</b> { <i>1</i>   <i>2</i> }]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# metric-style wide level 1	Configures a router to generate and accept only wide link metrics in the Level 1 area.
Step 5	<b>exit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# exit	Exits router address family configuration mode, and returns the router to router configuration mode.
Step 6	<b>interface</b> <i>type number</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3	Enters interface configuration mode.
Step 7	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } [ <b>unicast</b>   <b>multicast</b> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 8	<b>tag</b> <i>tag</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if-af)# tag 3	Sets the value of the tag to associate with the advertised connected route.

	Command or Action	Purpose
Step 9	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if-af)# end  OR  RP/0/RP0/CPU0:router(config-isis-if-af)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 10	<pre>show isis [ipv4   ipv6   afi-all] [unicast   multicast   safi-all] route [detail]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-isis-if-af)# show  isis ipv6 detail </p>	<p>Displays tag information. Verify that all tags are present in the RIB.</p>

## Setting the Priority for Adding Prefixes to the RIB

This optional task describes how to set the priority (order) for which specified prefixes are added to the RIB. The prefixes can be chosen using an access list (ACL), prefix list, or by matching a tag value.

### SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **address-family** { *ipv4* | *ipv6* } [ *unicast* | *multicast* ]
4. **metric-style wide** [ *transition* ] [ *level* { *1* | *2* } ]
5. **spf prefix-priority** [ *level* { *1* | *2* } ] { *critical* | *high* | *medium* } { *access-list-name* | *tag tag* }
6. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis instance-id</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. In this example, the IS-IS instance is called isp.
Step 3	<b>address-family {ipv4   ipv6} [unicast   multicast]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv6 address family.</li> </ul>
Step 4	<b>metric-style wide [transition] [level {1   2}]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# metric-style wide level 1	Configures a router to generate and accept only wide link metrics in the Level 1 area.
Step 5	<b>spf prefix-priority [level {1   2}] {critical   high   medium} {access-list-name   tag tag}</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority high tag 3	Installs all routes tagged with the value 3 first.
Step 6	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-af)# end OR RP/0/RP0/CPU0:router(config-isis-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring IP Fast Reroute Loop-free Alternate

This optional task describes how to enable the IP fast reroute (IPFRR) loop-free alternate (LFA) computation to converge traffic flows around link failures.



### Note

To enable node protection on broadcast links, IPFRR and bidirectional forwarding detection (BFD) must be enabled on the interface under IS-IS.

## Restrictions

IPFRR is supported on the Cisco CRS-1 router only. IPv4 address families and single-level interfaces are supported.

Multiprotocol Label Switching (MPLS) FRR and IPFRR cannot be configured on the same interface simultaneously.

## SUMMARY STEPS

1. **configure**
2. **router isis** *instance-id*
3. **interface** *type instance*
4. **circuit-type** {**level-1** | **level-2-only**}
5. **address-family ipv4** [**unicast**]
6. **ipfrr lfa** {**level** {**1** | **2**}}
7. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router isis</b> <i>instance-id</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. In this example, the IS-IS instance is called isp.
Step 3	<b>interface</b> <i>type instance</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis)# interface POS 0/1/0/3	Enters interface configuration mode.

	Command or Action	Purpose
Step 4	<b>circuit-type</b> {level-1   level-2-only}  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# circuit-type level-1	(Optional) Configures the type of adjacency.
Step 5	<b>address-family ipv4</b> [unicast]  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast	Specifies the IPv4 address family, and enters router address family configuration mode. <ul style="list-style-type: none"> <li>This example specifies the unicast IPv4 address family.</li> </ul>
Step 6	<b>ipfrr lfa</b> {level {1   2}}  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if-af)# ipfrr lfa level 1	Specifies the IP fast reroute loop-free alternate computation on link or node failures.
Step 7	<b>end</b> or <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-isis-if-af)# end or RP/0/RP0/CPU0:router(config-isis-if-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:   Uncommitted changes found, commit them before exiting (yes/no/cancel)?  [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuration Examples for Implementing IS-IS on Cisco IOS XR Software

This section provides the following configuration examples:

- [Configuring Single-Topology IS-IS for IPv6: Example, page RC-315](#)
- [Configuring Multitopology IS-IS for IPv6: Example, page RC-315](#)
- [Redistributing IS-IS Routes Between Multiple Instances: Example, page RC-315](#)
- [Tagging Routes: Example, page RC-316](#)



## Configuring Single-Topology IS-IS for IPv6: Example

The following example shows single-topology mode being enabled. An IS-IS instance is created, the NET is defined, IPv6 is configured along with IPv4 on an interface, and IPv4 link topology is used for IPv6.

This configuration allows POS interface 0/3/0/0 to form adjacencies for both IPv4 and IPv6 addresses.

```
router isis isp
  net 49.0000.0000.0001.00
  address-family ipv6 unicast
    single-topology
  interface POS0/3/0/0
    address-family ipv4 unicast
    !
    address-family ipv6 unicast
    !
  exit
!
interface POS0/3/0/0
  ipv4 address 10.0.1.3 255.255.255.0
  ipv6 address 2001::1/64
```

## Configuring Multitopology IS-IS for IPv6: Example

The following example shows multitopology IS-IS being configured in IPv6.

```
router isis isp
  net 49.0000.0000.0001.00
  interface POS0/3/0/0
    address-family ipv6 unicast
    metric-style wide level 1
  exit
!
interface POS0/3/0/0
  ipv6 address 2001::1/64
```

## Redistributing IS-IS Routes Between Multiple Instances: Example

The following example shows usage of the **set-attached-bit** and **redistribute** commands. Two instances, instance “1” restricted to Level 1 and instance “2” restricted to Level 2, are configured.

The Level 1 instance is propagating routes to the Level 2 instance using redistribution. Note that the administrative distance is explicitly configured higher on the Level 2 instance to ensure that Level 1 routes are preferred.

Attached bit is being set for the Level 1 instance since it is redistributing routes into the Level 2 instance. Therefore, instance “1” is a suitable candidate to get from the area to the backbone.

```
router isis 1
  is-type level-2-only
  net 49.0001.0001.0001.0001.00
  address-family ipv4 unicast
  distance 116
  redistribute isis 2 level 2
!
interface POS0/3/0/0
  address-family ipv4 unicast
!
```

```

!
router isis 2
 is-type level-1
 net 49.0002.0001.0001.0002.00
 address-family ipv4 unicast
  set-attached-bit
!
interface POS0/1/0/0
 address-family ipv4 unicast

```

## Tagging Routes: Example

The following example show how to tag routes.

```

route-policy isis-tag-55
end-policy
!
route-policy isis-tag-555
 if destination in (5.5.5.0/24 eq 24) then
  set tag 555
  pass
 else
  drop
 endif
end-policy
!
router static
 address-family ipv4 unicast
  0.0.0.0/0 2.6.0.1
  5.5.5.0/24 Null0
!
!
router isis uut
 net 00.0000.0000.12a5.00
 address-family ipv4 unicast
 metric-style wide
 redistribute static level-1 route-policy isis-tag-555
 spf prefix-priority critical tag 13
 spf prefix-priority high tag 444
 spf prefix-priority medium tag 777

```

## Where to Go Next

To implement more IP routing protocols, see the following document modules in *Cisco IOS XR Routing Configuration Guide*:

- *Implementing OSPF on Cisco IOS XR Software*
- *Implementing BGP on Cisco IOS XR Software*
- *Implementing EIGRP on Cisco IOS XR Software*
- *Implementing RIP on Cisco IOS XR Software*

# Additional References

The following sections provide references related to implementing IS-IS on Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
IS-IS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference</i> , Release 3.6
MPLS TE feature information	<i>Implementing MPLS Traffic Engineering on Cisco IOS XR Software</i> module in the <i>Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i> , Release 3.6
IS-IS TLVs	Intermediate System-to-Intermediate System (IS-IS) TLVs at: <a href="http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094bbd.shtml">http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094bbd.shtml</a>
Bidirectional Forwarding Detection (BFD)	<i>Cisco IOS XR Interface and Hardware Configuration Guide</i> , Release 3.5 and <i>Cisco IOS XR Interface and Hardware Command Reference</i> , Release 3.6

## Standards

Standards	Title
Draft-ietf-isis-ipv6-05.txt	<i>Routing IPv6 with IS-IS</i> , by Christian E. Hopps
Draft-ietf-isis-wg-multi-topology-06.txt	<i>M-ISIS: Multi Topology (MT) Routing in IS-IS</i> , by Tony Przygienda, Naiming Shen, and Nischal Sheth
Draft-ietf-isis-traffic-05.txt	<i>IS-IS Extensions for Traffic Engineering</i> , by Henk Smit and Toni Li
Draft-ietf-isis-restart-04.txt	<i>Restart Signalling for IS-IS</i> , by M. Shand and Les Ginsberg
Draft-ietf-isis-igp-p2p-over-lan-05.txt	<i>Point-to-point operation over LAN in link-state routing protocols</i> , by Naiming Shen
Draft-ietf-rtgwg-ipfr-framework-06.txt	<i>IP Fast Reroute Framework</i> , by M. Shand and S. Bryant
Draft-ietf-rtgwg-lf-conv-frmwk-00.txt	<i>A Framework for Loop-free Convergence</i> , by M. Shand and S. Bryant

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
RFC 1142	<i>OSI IS-IS Intra-domain Routing Protocol</i>
RFC 1195	<i>Use of OSI IS-IS for Routing in TCP/IP and Dual Environments</i>
RFC 2763	<i>Dynamic Hostname Exchange Mechanism for IS-IS</i>
RFC 2966	<i>Domain-wide Prefix Distribution with Two-Level IS-IS</i>
RFC 2973	<i>IS-IS Mesh Groups</i>
RFC 3277	<i>IS-IS Transient Blackhole Avoidance</i>
RFC 3373	<i>Three-Way Handshake for IS-IS Point-to-Point Adjacencies</i>
RFC 3567	<i>IS-IS Cryptographic Authentication</i>

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>



# Implementing and Monitoring RIB on Cisco IOS XR Software

Routing Information Base (RIB) is a distributed collection of information about routing connectivity among all nodes of a network.

Each router maintains a RIB containing the routing information for that router. RIB stores the best routes from all routing protocols that are running on the system.

This module describes the tasks you need to perform to implement and monitor RIB on your Cisco IOS XR network.



## Note

For more information about RIB on the Cisco IOS XR software and complete descriptions of RIB commands listed in this module, see the [“Related Documents”](#) of this module. To locate documentation for other commands that might appear during the execution of a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing and Monitoring RIB on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router.
Release 3.3.0	VPN routing and forwarding (VRF) support was added to the command syntax.
Release 3.4.0	RIB statistics support was added using the <b>show rib statistics</b> command. Disabling RIB next-hop dampening was supported.
Release 3.5.0	The following features were supported: <ul style="list-style-type: none"><li>• IP fast reroute loop-free alternates computation</li><li>• IPv6 Provider Edge and IPv6 VPN Provider Edge over Multiprotocol Label Switching</li><li>• RIB quarantining</li></ul>
Release 3.6.0	No modification.

# Contents

- [Prerequisites for Implementing RIB on Cisco IOS XR Software, page RC-320](#)
- [Information About RIB Configuration, page RC-320](#)
- [How to Deploy and Monitor RIB, page RC-324](#)
- [Configuration Examples for RIB Monitoring, page RC-327](#)
- [Where to Go Next, page RC-330](#)
- [Additional References, page RC-330](#)

## Prerequisites for Implementing RIB on Cisco IOS XR Software

- To use this command, you must be in a user group associated with a task group that includes the proper task IDs. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of the *Cisco IOS XR System Security Configuration Guide*.
- RIB is distributed with the base Cisco IOS XR software; as such, it does not have any special requirements for installation. The following are the requirements for base software installation:
  - Router
  - Cisco IOS XR software
  - Base package

## Information About RIB Configuration

To implement the Cisco RIB feature, you must understand the following concepts:

- [Overview of RIB, page RC-320](#)
- [RIB Data Structures in BGP and Other Protocols, page RC-321](#)
- [RIB Administrative Distance, page RC-321](#)
- [RIB Support for IPv4 and IPv6, page RC-322](#)
- [RIB Statistics, page RC-322](#)
- [IPv6 and IPv6 VPN Provider Edge Transport over MPLS, page RC-322](#)
- [IP Fast Reroute, page RC-323](#)
- [RIB Quarantining, page RC-323](#)

## Overview of RIB

Each routing protocol selects its own set of best routes and installs those routes and their attributes in RIB. RIB stores these routes and selects the best ones from among all routing protocols. Those routes are downloaded to the line cards for use in forwarding packets. The acronym RIB is used both to refer to RIB processes and the collection of route data contained within RIB.

Within a protocol, routes are selected based on the metrics in use by that protocol. A protocol downloads its best routes (lowest or tied metric) to RIB. RIB selects the best overall route by comparing the administrative distance of the associated protocol.

## RIB Data Structures in BGP and Other Protocols

RIB uses processes and maintains data structures distinct from other routing applications, such as Border Gateway Protocol (BGP) and other unicast routing protocols, or multicast protocols, such as Protocol Independent Multicast (PIM) or Multicast Source Discovery Protocol (MSDP). However, these routing protocols use internal data structures similar to what RIB uses, and may internally refer to the data structures as a RIB. For example, BGP routes are stored in the BGP RIB (BRIB), and multicast routes, computed by multicast routing protocols such as PIM and MSDP, are stored in the Multicast RIB (MRIB). RIB processes are not responsible for the BRIB and MRIB, which are handled by BGP and multicast processes, respectively.

The table used by the line cards and RP to forward packets is called the Forwarding Information Base (FIB). RIB processes do not build the FIBs. Instead, RIB downloads the set of selected best routes to the FIB processes, by the Bulk Content Downloader (BCDL) process, onto each line card. FIBs are then constructed.

## RIB Administrative Distance

Forwarding is done based on the longest prefix match. If you are forwarding a packet destined to 10.0.2.1, you prefer 10.0.2.0/24 over 10.0.0.0/16 because the mask /24 is longer (and more specific) than a /16.

Routes from different protocols that have the same prefix and length are chosen based on administrative distance. For instance, the Open Shortest Path First (OSPF) protocol has an administrative distance of 110, and the Intermediate System-to-Intermediate System (IS-IS) protocol has an administrative distance of 115. If IS-IS and OSPF both download 10.0.1.0/24 to RIB, RIB would prefer the OSPF route because OSPF has a lower administrative distance. Administrative distance is used only to choose between multiple routes of the same length.

The default administrative distances for the common protocols are shown in [Table 1](#).

**Table 1**      **Default Administrative Distances**

Protocol	Administrative Distance Default
Connected or local routes	0
Static routes	1
External BGP routes	20
OSPF routes	110
IS-IS routes	115
Internal BGP routes	200

The administrative distance for some routing protocols (for instance IS-IS, OSPF, and BGP) can be changed. See the protocol-specific documentation for the proper method to change the administrative distance of that protocol.

**Note**

Changing the administrative distance of a protocol on some but not all routers can lead to routing loops and other undesirable behavior. Doing so is not recommended.

## RIB Support for IPv4 and IPv6

In Cisco IOS XR software, RIB tables support multicast and unicast routing.

The default routing table for Cisco IOS XR RIB are the unicast and the multicast-unicast RIB tables for IPv4 and IPv6 routing, respectively. For multicast routing, routing protocols insert unicast routes into the multicast-unicast RIB table. Multicast protocols then use the information to build multicast routes (which in turn are stored in the MRIB). See the multicast documentation for more information on using and configuring multicast.

RIB processes `ipv4_rib` and `ipv6_rib` run on the RP card. If process placement functionality is available and supported by multiple RPs in the router, RIB processes can be placed on any available node.

## RIB Statistics

RIB supports statistics for messages (requests) flowing between the RIB and its clients. Protocol clients send messages to the RIB (for example, route add, route delete, and next-hop register, and so on). RIB also sends messages (for example, redistribute routes, advertisements, next-hop notifications, and so on). These statistics are used to gather information about what messages have been sent and the number of messages that have been sent. These statistics provide counters for the various messages that flow between the RIB server and its clients. The statistics are displayed using the **show rib statistics** command.

RIB maintains counters for all requests sent from a client including:

- Route operations
- Table registrations
- Next-hop registrations
- Redistribution registrations
- Attribute registrations
- Synchronization completion

RIB also maintains counters for all requests sent by the RIB. The configuration will disable the RIB next-hop dampening feature. As a result, RIB notifies client immediately when a next hop that client registered for is resolved or unresolved.

RIB also maintains the results of the requests.

## IPv6 and IPv6 VPN Provider Edge Transport over MPLS

IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) leverages the existing Multiprotocol Label Switching (MPLS) IPv4 core infrastructure for IPv6 transport. 6PE and 6VPE enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs).



RIB supports 6PE and 6VPE by providing 6VPE next hops. The next-hop information is stored in an opaque database in RIB, which is populated by protocol clients with data to be sent to the Forwarding Information Base (FIB).

For detailed information about configuring 6PE and 6VPE over MPLS, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

## IP Fast Reroute

The IP Fast Reroute (IPFRR) loop-free alternate (LFA) computation provides protection against link failure. Locally computed repair paths are used to prevent packet loss caused by loops that occur during network reconvergence after a failure. For information about IPFRR see *Implementing IS-IS on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*.



**Note**

---

IPFRR is supported on the Cisco CRS-1 router.

---

## RIB Quarantining

RIB quarantining solves the problem in the interaction between routing protocols and the RIB. The problem is a persistent oscillation between the RIB and routing protocols that occurs when a route is continuously inserted and then withdrawn from the RIB, resulting in a spike in CPU use until the problem is resolved. If there is no damping on the oscillation, then both the protocol process and the RIB process have high CPU use, affecting the rest of the system as well as blocking out other protocol and RIB operations. This problem occurs when a particular combination of routes is received and installed in the RIB. This problem typically happens as a result of a network misconfiguration. However, because the misconfiguration is across the network, it is not possible to detect the problem at configuration time on any single router.

The quarantining mechanism detects mutually recursive routes and quarantines the last route that completes the mutual recursion. The quarantined route is periodically evaluated to see if the mutual recursion has gone away. If the recursion still exists, the route remains quarantined. If the recursion has gone away, the route is released from its quarantine.

The following steps are used to quarantine a route:

1. RIB detects when a particular problematic path is installed.
2. RIB sends a notification to the protocol that installed the path.
3. When the protocol receives the quarantine notification about the problem route, it marks the route as being “quarantined.” If it is a BGP route, BGP does not advertise reachability for the route to its neighbors.
4. Periodically, RIB tests all its quarantined paths to see if they can now safely be installed (moved from quarantined to “Ok to use” state). A notification is sent to the protocol to indicate that the path is now safe to use.

# How to Deploy and Monitor RIB

To deploy and monitor RIB, you must understand the following concepts:

- [Verifying RIB Configuration Using the Routing Table, page RC-324](#) (required)
- [Verifying Networking and Routing Problems, page RC-324](#) (required)
- [Disabling RIB Next-hop Dampening, page RC-326](#) (optional)

## Verifying RIB Configuration Using the Routing Table

Perform this task to verify the RIB configuration to ensure that RIB is running on the RP and functioning properly by checking the routing table summary and details.

### SUMMARY STEPS

1. **show route** [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **summary** [**detail**] [**standby**]
2. **show route** [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] [*protocol* [*instance*] | *ip-address* [*mask*] | *ip-address/prefix-length*] [**standby**] [**detail**]

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] summary [detail] [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route summary</p>	Displays route summary information about the specified routing table. <ul style="list-style-type: none"> <li>• The default table summarized is the IPv4 unicast routing table.</li> </ul>
Step 2	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] [protocol [instance]   ip-address mask] [standby] [detail]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast</p>	Displays more detailed route information about the specified routing table. <ul style="list-style-type: none"> <li>• This command is usually issued with an IP address or other optional filters to limit its display. Otherwise, it displays all routes from the default IPv4 unicast routing table, which can result in an extensive list, depending on the configuration of the network.</li> </ul>

## Verifying Networking and Routing Problems

Perform this task to verify the operation of routes between nodes.

### SUMMARY STEPS

1. **show route** [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] [*protocol* [*instance*] | *ip-address mask*] [**standby**] [**detail**]

2. `show route [vrf {vrf-name | all}] [afi-all | ipv4 | ipv6] [unicast | multicast | safi-all] backup [ip-address] [standby]`
3. `show route [vrf {vrf-name | all}] [ipv4 | ipv6] [unicast | multicast | safi-all] best-local ip-address [standby]`
4. `show route [vrf {vrf-name | all}] [afi-all | ipv4 | ipv6] [unicast | multicast | safi-all] connected [standby]`
5. `show route [vrf {vrf-name | all}] [afi-all | ipv4 | ipv6] [unicast | multicast | safi-all] local [interface] [standby]`
6. `show route [vrf {vrf-name | all}] [ipv4 | ipv6] [unicast | multicast | safi-all] longer-prefixes {ip-address mask | ip-address/prefix-length} [standby]`
7. `show route [vrf {vrf-name | all}] [ipv4 | ipv6] [unicast | multicast | safi-all] next-hop [ip-address] [standby]`

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] [protocol [instance]   ip-address mask] [standby] [detail]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route list list1 bgp aspo ipv4 unicast 192.168.111/8</p>	Displays the current routes in RIB.
Step 2	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] backup [ip-address] [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast backup 192.168.111/8</p>	Displays backup routes in RIB.
Step 3	<pre>show route [vrf {vrf-name   all}] [ipv4   ipv6] [unicast   multicast   safi-all] best-local ip-address [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast best-local 192.168.111/8</p>	Displays the best-local address to use for return packets from the given destination.
Step 4	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] connected [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast connected</p>	Displays the current connected routes of the routing table.

	Command or Action	Purpose
Step 5	<pre>show route [vrf {vrf-name   all}] [afi-all   ipv4   ipv6] [unicast   multicast   safi-all] local [interface] [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast local</p>	Displays local routes for receive entries in the routing table.
Step 6	<pre>show route [vrf {vrf-name   all}] [ipv4   ipv6] [unicast   multicast   safi-all] longer-prefixes {ip-address mask   ip-address/prefix-length} [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast 192.168.111/8 longer-prefixes</p>	Displays the current routes in RIB that share a given number of bits with a given network.
Step 7	<pre>show route [vrf {vrf-name   all}] [ipv4   ipv6] [unicast   multicast   safi-all] next-hop ip-address [standby]</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# show route ipv4 unicast next-hop 192.168.1.34</p>	Displays the next-hop gateway or host to a destination address.

## Disabling RIB Next-hop Dampening

Perform this task to disable RIB next-hop dampening.

### SUMMARY STEPS

1. **router rib**
2. **address-family {ipv4 | ipv6} next-hop dampening disable**
3. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>router rib</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router# route rib</p>	Enters RIB configuration mode.
Step 2	<pre>address-family {ipv4   ipv6} next-hop dampening disable</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rib)# address family ipv4 next-hop dampening disable</p>	Disables next-hop dampening for IPv4 address families.
Step 3	<pre>end</pre> <p>OR</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rib)# end OR RP/0/RP0/CPU0:router(config-rib)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuration Examples for RIB Monitoring

RIB is not configured separately for the Cisco IOS XR system. RIB computes connectivity of the router with other nodes in the network based on input from the routing protocols. RIB may be used to monitor and troubleshoot the connections between RIB and its clients, but it is essentially used to monitor routing connectivity between the nodes in a network. This section contains displays from the **show** commands used to monitor that activity. The following sample output is provided:

- [Output of show route Command: Example, page RC-328](#)
- [Output of show route backup Command: Example, page RC-328](#)
- [Output of show route best-local Command: Example, page RC-328](#)
- [Output of show route connected Command: Example, page RC-329](#)
- [Output of show route local Command: Example, page RC-329](#)

- [Output of show route longer-prefixes Command: Example, page RC-329](#)
- [Output of show route next-hop Command: Example, page RC-329](#)

## Output of show route Command: Example

The following is sample output from the **show route** command when entered without an address:

```
RP/0/RP0/CPU0:router# show route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

Gateway of last resort is 172.23.54.1 to network 0.0.0.0

C    10.2.210.0/24 is directly connected, 1d21h, Ethernet0/1/0/0
L    10.2.210.221/32 is directly connected, 1d21h, Ethernet0/1/1/0
C    172.20.16.0/24 is directly connected, 1d21h, ATM4/0.1
L    172.20.16.1/32 is directly connected, 1d21h, ATM4/0.1
C    10.6.100.0/24 is directly connected, 1d21h, Loopback1
L    10.6.200.21/32 is directly connected, 1d21h, Loopback0
S    192.168.40.0/24 [1/0] via 172.20.16.6, 1d21h
```

## Output of show route backup Command: Example

The following is sample output from the **show route backup** command:

```
RP/0/RP0/CPU0:router# show route backup

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

S    172.73.51.0/24 is directly connected, 2d20h, POS4/0/0/1
      Backup O E2 [110/1] via 10.12.12.2, POS3/0/0/1
```

## Output of show route best-local Command: Example

The following is sample output from the **show route best-local** command:

```
RP/0/RP0/CPU0:router# show route best-local 10.12.12.1

Routing entry for 10.12.12.1/32
  Known via "local", distance 0, metric 0 (connected)
  Routing Descriptor Blocks
    10.12.12.1 directly connected, via POS3/0
    Route metric is 0
```

## Output of show route connected Command: Example

The following is sample output from the **show route connected** command:

```
RP/0/RP0/CPU0:router# show route connected

C    10.2.210.0/24 is directly connected, 1d21h, Ethernet0
C    172.20.16.0/24 is directly connected, 1d21h, ATM4/0.1
C    10.6.100.0/24 is directly connected, 1d21h, Loopback1
```

## Output of show route local Command: Example

The following is sample output from the **show route local** command:

```
RP/0/RP0/CPU0:router# show route local

L    10.10.10.1/32 is directly connected, 00:14:36, Loopback0
L    10.91.36.98/32 is directly connected, 00:14:32, Ethernet0/0
L    172.22.12.1/32 is directly connected, 00:13:35, POS3/0
L    192.168.20.2/32 is directly connected, 00:13:27, GigabitEthernet2/0
L    10.254.254.1/32 is directly connected, 00:13:26, GigabitEthernet2/2
```

## Output of show route longer-prefixes Command: Example

The following is sample output from the **show route longer-prefixes** command:

```
RP/0/RP0/CPU0:router# show route ipv4 172.16.0.0/8 longer-prefixes

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1
       E2 - OSPF external type 2, E - EGP, i - ISIS, L1 - IS-IS level-1
       L2 - IS-IS level-2, ia - IS-IS inter area
       su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

Gateway of last resort is 172.23.54.1 to network 0.0.0.0
S    172.16.2.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.3.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.4.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.5.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.6.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.7.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.8.0/32 is directly connected, 00:00:24, Loopback0
S    172.16.9.0/32 is directly connected, 00:00:24, Loopback0
```

## Output of show route next-hop Command: Example

The following is sample output from the **show route resolving-next-hop** command:

```
RP/0/RP0/CPU0:router# show route resolving-next-hop 10.0.0.1

Nexthop matches 0.0.0.0/0
  Known via "static", distance 200, metric 0, candidate default path
  Installed Aug 18 00:59:04.448
  Directly connected nexthops
    172.29.52.1, via MgmtEth0/RP0/CPU0/0
    Route metric is 0
```

```
172.29.52.1, via MgmtEth0/RP1/CPU0/0  
Route metric is 0
```

## Where to Go Next

For additional information on the protocols that interact with RIB, you may want to see the following publications:

- *Implementing MPLS Layer 3 VPNs in Cisco IOS XR Multiprotocol Label Switching Configuration Guide.*
- *Implementing BGP on Cisco IOS XR Software in Cisco IOS XR Routing Configuration Guide.*
- *Implementing EIGRP on Cisco IOS XR Software in Cisco IOS XR Routing Configuration Guide.*
- *Implementing IS-IS on Cisco IOS XR Software in Cisco IOS XR Routing Configuration Guide.*
- *Implementing OSPF on Cisco IOS XR Software in Cisco IOS XR Routing Configuration Guide.*
- *Implementing RIP on Cisco IOS XR Software in Cisco IOS XR Routing Configuration Guide.*
- *RIB Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference.*

## Additional References

The following sections provide references related to implementing RIB on Cisco IOS XR software:



## Related Documents

Related Topic	Document Title
Routing Information Base commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>RIB Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
BGP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>BGP Commands on Cisco IOS XR Software, in Cisco IOS XR Routing Command Reference</i>
IS-IS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>IS-IS Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
OSPF commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>OSPF Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
OSPFv3 commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>OSPFv3 Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
EIGRP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>EIGRP Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
RIP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>RIP Commands on Cisco IOS XR Software in Cisco IOS XR Routing Command Reference</i>
Multicast commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Multicast Command Reference</i>
Multicast configuration: configuration concepts, task, and examples	<i>Cisco IOS XR Multicast Configuration Guide</i>
MPLS Layer 3 VPN configuration: configuration concepts, task, and examples	<i>Implementing MPLS Layer 3 VPNs in Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>

## Standards

Standards	Title
Draft-ietf-rtgwg-ipfrr-framework-06.txt	<i>IP Fast Reroute Framework</i> , by M. Shand and S. Bryant
Draft-ietf-rtgwg-lf-conv-frmwk-00.txt	<i>A Framework for Loop-free Convergence</i> , by M. Shand and S. Bryant

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>



# Implementing RIP on Cisco IOS XR Software

The Routing Information Protocol (RIP) is a classic distance vector Interior Gateway Protocol (IGP) designed to exchange information within an autonomous system (AS) of a small network.

This module describes the concepts and tasks you need to implement basic RIP routing on your router. Cisco IOS XR software supports a standard implementation of RIP Version 2 (RIPv2) that supports backward compatibility with RIP Version 1 (RIPv1) as specified by RFC2453.

For RIP configuration information related to the following features, see the “[Related Documents](#)” section of this module.

- Multiprotocol Label Switching (MPLS) Layer 3 Virtual Private Network (VPN)
- Site of Origin (SoO) Support



## Note

For more information about RIP on the Cisco IOS XR software and complete descriptions of the RIP commands listed in this module, see the “[Related Documents](#)” section of this module. To locate documentation for other commands that might appear while performing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing RIP on Cisco IOS XR Software

Release	Modification
Release 3.3.0	This feature was introduced on the Cisco CRS-1 and Cisco XR 12000 Series Router.
Release 3.4.0	No modification.
Release 3.5.0	Four-byte autonomous system (AS) number support was added.
Release 3.6.0	No modification.

## Contents

- [Information About Implementing RIP on Cisco IOS XR Software](#), page RC-334
- [How to Implement RIP on Cisco IOS XR Software](#), page RC-337
- [Configuration Examples for Implementing RIP on Cisco IOS XR Software](#), page RC-347
- [Additional References](#), page RC-350

# Information About Implementing RIP on Cisco IOS XR Software

- [Prerequisites for Implementing RIP on Cisco IOS XR Software, page RC-334](#)
- [RIP Functional Overview, page RC-334](#)
- [Split Horizon for RIP, page RC-335](#)
- [Route Timers for RIP, page RC-335](#)
- [Route Redistribution for RIP, page RC-336](#)
- [Default Administrative Distances for RIP, page RC-336](#)
- [Routing Policy Options for RIP, page RC-337](#)

## Prerequisites for Implementing RIP on Cisco IOS XR Software

The following prerequisite must be met for implementing RIP on Cisco IOS XR software. You must be in a user group associated with a task group that includes the proper task IDs for RIP commands. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of the *Cisco IOS XR System Security Configuration Guide*.

## RIP Functional Overview

RIP Version 1 (RIP v1) is a classful, distance-vector protocol that is considered the easiest routing protocol to implement. Unlike OSPF, RIP broadcasts User Datagram Protocol (UDP) data packets to exchange routing information in internetworks that are flat rather than hierarchical. Network complexity and network management time is reduced. However, as a classful routing protocol, RIP v1 allows only contiguous blocks of hosts, subnets or networks to be represented by a single route, severely limiting its usefulness.

RIP v2 allows more information carried in RIP update packets, such as support for:

- Route summarization
- Classless interdomain routing (CIDR)
- Variable-length subnet masks (VLSMs)
- Autonomous systems and the use of redistribution
- Multicast address 224.0.0.9 for RIP advertisements

The metric that RIP uses to rate the value of different routes is *hop count*. The hop count is the number of routers that can be traversed in a route. A directly connected network has a metric of zero; an unreachable network has a metric of 16. This small range of metrics makes RIP an unsuitable routing protocol for large networks.

Routing information updates are advertised every 30 seconds by default, and new updates discovered from neighbor routers are stored in a routing table.

Only RIP Version 2 (RIP v2), as specified in RFC 2453, is supported on Cisco IOS XR software and, by default, the software only sends and receives RIP v2 packets. However, you can configure the software to send, or receive, or both, only Version 1 packets or only Version 2 packets or both version type packets per interface.

Here are some good reasons to use RIP:

- Compatible with diverse network devices
- Best for small networks, because there is very little overhead, in terms of bandwidth used, configuration, and management time.
- Support for legacy host systems

Because of RIP's ease of use, it is implemented in networks worldwide.

## Split Horizon for RIP

Normally, routers that are connected to broadcast-type IP networks and that use distance-vector routing protocols employ the *split horizon* mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routers, particularly when links are broken.

If an interface is configured with secondary IP addresses and split horizon is enabled, updates might not be sourced by every secondary address. One routing update is sourced per network number unless split horizon is disabled.



### Note

---

The split horizon feature is enabled by default. In general, we recommend that you do not change the default state of split horizon unless you are certain that your operation requires the change in order to properly advertise routes.

---

## Route Timers for RIP

RIP uses several timers that determine such variables as the frequency of routing updates, the length of time before a route becomes invalid, and other parameters. You can adjust these timers to tune routing protocol performance to better suit your internetwork needs, by making the following timer adjustments to:

- The rate (time in seconds between updates) at which routing updates are sent
- The interval of time (in seconds) after which a route is declared invalid
- The interval (in seconds) during which routing information regarding better paths is suppressed
- The amount of time (in seconds) that must pass before a route is removed from the RIP topology table
- The amount of time delay between RIP update packets

The first four timer adjustments are configurable by the **timers basic** command. The **output-delay** command changes the amount of time delay between RIP update packets. See the [“Customize RIP” section on page RC-339](#) for configuration details.

It also is possible to tune the IP routing support in the software to enable faster convergence of the various IP routing algorithms and quickly drop back to redundant routers, if necessary. The total result is to minimize disruptions to end users of the network in situations in which quick recovery is essential.

## Route Redistribution for RIP

Redistribution is a feature that allows different routing domains, to exchange routing information. Networking devices that route between different routing domains are called *boundary routers*, and it is these devices that inject the routes from one routing protocol into another. Routers within a routing domain only have knowledge of routes internal to the domain unless route redistribution is implemented on the *boundary routers*.

When running RIP in your routing domain, you might find it necessary to use multiple routing protocols within your internetwork and redistribute routes between them. Some common reasons are:

- To advertise routes from other protocols into RIP, such as static, connected, OSPF and BGP.
- To migrate from RIP to a new Interior Gateway Protocol (IGP) such as EIGRP.
- To retain routing protocol on some routers to support host systems, but upgrade routers for other department groups.
- To communicate among a mixed-router vendor environment. Basically, you might use a protocol specific to Cisco in one portion of your network and use RIP to communicate with devices other than Cisco devices.

Further, route redistribution gives a company the ability to run different routing protocols in work groups or areas in which each is particularly effective. By not restricting customers to using only a single routing protocol, Cisco IOS XR route redistribution is a powerful feature that minimizes cost, while maximizing technical advantage through diversity.

When it comes to implementing route redistribution in your internetwork, it can be very simple or very complex. An example of a simple one-way redistribution is to log into a router on which RIP is enabled and use the **redistribute static** command to advertise only the static connections to the backbone network to pass through the RIP network. For complex cases in which you must consider routing loops, incompatible routing information, and inconsistent convergence time, you must determine why these problems occur by examining how Cisco routers select the best path when more than one routing protocol is running administrative cost.

## Default Administrative Distances for RIP

Administrative distance is used as a measure of the trustworthiness of the source of the IP routing information. When a dynamic routing protocol such as RIP is configured, and you want to use the redistribution feature to exchange routing information, it is important to know the default administrative distances for other route sources so that you can set the appropriate distance weight. See [Table 1](#) for a list of the default administrative distances.

**Table 1**      **Default Administrative Distances of Routing Protocols**

Routing Protocols	Administrative Distance Value
Connected interface	0
Static route out an interface	0
Static route to next hop	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110

**Table 1**      **Default Administrative Distances of Routing Protocols (continued)**

Routing Protocols	Administrative Distance Value
IS-IS	115
RIP version 1 and 2	120
External EIGRP	170
Internal BGP	200
Unknown	255

An administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored. Administrative distance values are subjective; there is no quantitative method for choosing them.

## Routing Policy Options for RIP

Route policies comprise series of statements and expressions that are bracketed with the **route-policy** and **end-policy** keywords. Rather than a collection of individual commands (one for each line), the statements within a route policy have context relative to each other. Thus, instead of each line being an individual command, each policy or set is an independent configuration object that can be used, entered, and manipulated as a unit.

Each line of a policy configuration is a logical subunit. At least one new line must follow the **then**, **else**, and **end-policy** keywords. A new line must also follow the closing parenthesis of a parameter list and the name string in a reference to an AS path set, community set, extended community set, or prefix set. At least one new line must precede the definition of a route policy, AS path set, community set, extended community set, or prefix set. One or more new lines can follow an action statement. One or more new lines can follow a comma separator in a named AS path set, community set, extended community set, or prefix set. A new line must appear at the end of a logical unit of policy expression and may not appear anywhere else.

## How to Implement RIP on Cisco IOS XR Software

This section contains instructions for the following tasks:

- [Enabling RIP, page RC-338](#) (required)
- [Customize RIP, page RC-339](#) (optional)
- [Control Routing Information, page RC-342](#) (optional)
- [Creating a Route Policy for RIP, page RC-344](#) (optional)

**Note**

To save configuration changes, you must commit changes when the system prompts you.

## Enabling RIP

This task enables RIP routing and establishes a RIP routing process.

### Prerequisites

Although you can configure RIP before you configure an IP address, no RIP routing occurs until at least one IP address is configured.

### SUMMARY STEPS

1. **configure**
2. **router rip**
3. **neighbor** *ip-address*
4. **broadcast-for-v2**
5. **interface** *type instance*
6. **receive version** {1 | 2 | 1 2}
7. **send version** {1 | 2 | 1 2}
8. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router rip</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router rip	Configures a RIP routing process.
Step 3	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# neighbor 172.160.1.2	(Optional) Defines a neighboring router with which to exchange RIP protocol information.
Step 4	<b>broadcast-for-v2</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# broadcast-for-v2	(Optional) Configures RIP to send only Version 2 packets to the broadcast IP address rather than the RIP v2 multicast address (224.0.0.9). This command can be applied at the interface or global configuration level.



	Command or Action	Purpose
Step 5	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# interface pos 0/1/0/0</p>	(Optional) Defines the interfaces on which the RIP routing protocol runs.
Step 6	<p><b>receive version</b> {1   2   1 2}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip-if)# receive version 1 2</p>	(Optional) Configures an interface to accept packets that are: <ul style="list-style-type: none"> <li>• Only RIP v1</li> <li>• Only RIP v2</li> <li>• Both RIP v1 and RIP v2</li> </ul>
Step 7	<p><b>send version</b> {1   2   1 2}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip-if)# send version 1 2</p>	(Optional) Configures an interface to send packets that are: <ul style="list-style-type: none"> <li>• Only RIP v1</li> <li>• Only RIP v2</li> <li>• Both RIP v1 and RIP v2</li> </ul>
Step 8	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip-if)# end or RP/0/RP0/CPU0:router(config-rip-if)# commit</p>	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:               <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Customize RIP

This task describes how to customize RIP for network timing and the acceptance of route entries.

### SUMMARY STEPS

1. **configure**
2. **router rip**
3. **auto-summary**
4. **timers basic update invalid holddown flush**

5. **output-delay** *delay*
6. **nsf**
7. **interface** *type instance*
8. **metric-zero-accept**
9. **split-horizon** **disable**
10. **poison-reverse**
11. **end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router rip</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router rip	Configures a RIP routing process.
Step 3	<b>auto-summary</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# auto-summary	(Optional) Enables automatic route summarization of subnet routes into network-level routes. <ul style="list-style-type: none"> <li>• By default, auto-summary is disabled.</li> </ul> <b>Note</b> If you have disconnected subnets, use the <b>no</b> keyword to disable automatic route summarization and permit software to send subnet and host routing information across classful network boundaries.
Step 4	<b>timers basic</b> <i>update invalid holddown flush</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# timers basic 5 15 15 30	(Optional) Adjusts RIP network timers.  <b>Note</b> To view the current and default timer values, view output from the <b>show rip</b> command.
Step 5	<b>output-delay</b> <i>delay</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# output-delay 10	(Optional) Changes the interpacket delay for the RIP updates sent.  <b>Note</b> Use this command if you have a high-end router sending at high speed to a low-speed router that might not be able to receive at that fast a rate.
Step 6	<b>nsf</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# nsf	(Optional) Configures NSF on RIP routes after a RIP process shutdown or restart.

	Command or Action	Purpose
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip)# interface pos  0/1/0/0</p>	(Optional) Defines the interfaces on which the RIP routing protocol runs.
Step 8	<p><b>metric-zero-accept</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)#  metric-zero-accept</p>	(Optional) Allows the networking device to accept route entries received in update packets with a metric of zero (0). The received route entry is set to a metric of one (1).
Step 9	<p><b>split-horizon disable</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)#  split-horizon disable</p>	<p>(Optional) Disables the split horizon mechanism.</p> <ul style="list-style-type: none"> <li>• By default, split horizon is enabled.</li> <li>• In general, we do not recommend changing the state of the default for the <b>split-horizon</b> command, unless you are certain that your application requires a change to properly advertise routes. If split horizon is disabled on a serial interface (and that interface is attached to a packet-switched network), you must disable split horizon for all networking devices in any relevant multicast groups on that network.</li> </ul>
Step 10	<p><b>poison-reverse</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)#  poison-reverse</p>	Enables poison reverse processing of RIP router updates.
Step 11	<p><b>end</b>  or  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)# end  OR  RP/0/RP0/CPU0:router(config-rip-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>– Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>– Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>– Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Control Routing Information

This task describes how to control or prevent routing update exchange and propagation.

Some reasons to control or prevent routing updates are:

- To slow or stop the update traffic on a WAN link—If you do not control update traffic on an on-demand WAN link, the link remains up constantly. By default, RIP routing updates occur every 30 seconds.
- To prevent routing loops—If you have redundant paths or are redistributing routes into another routing domain, you may want to filter the propagation of one of the paths.
- To filter network received in updates — If you do not want other routers from learning a particular device's interpretation of one or more routes, you can suppress that information.
- To prevent other routers from processing routes dynamically— If you do not want to process routing updates entering the interface, you can suppress that information.
- To preserve bandwidth—You can ensure maximum bandwidth availability for data traffic by reducing unnecessary routing update traffic.

### SUMMARY STEPS

1. **configure**
2. **router rip**
3. **neighbor** *ip-address*
4. **interface** *type instance*
5. **passive-interface**
6. **exit**
7. **interface** *type instance*
8. **route-policy** {**in** | **out**}
9. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router rip</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router rip	Configures a RIP routing process.

	Command or Action	Purpose
Step 3	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip)# neighbor  172.160.1.2</p>	(Optional) Defines a neighboring router with which to exchange RIP protocol information.
Step 4	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip)# interface pos  0/1/0/0</p>	(Optional) Defines the interfaces on which the RIP routing protocol runs.
Step 5	<p><b>passive-interface</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)#  passive-interface</p>	(Optional) Suppresses the sending of RIP updates on an interface, but not to explicitly configured neighbors.
Step 6	<p><b>exit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip-if)# exit</p>	(Optional) Returns the router to the next higher configuration mode.
Step 7	<p><b>interface</b> <i>type instance</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rip)# interface pos  0/2/0/0</p>	(Optional) Defines the interfaces on which the RIP routing protocol runs.

	Command or Action	Purpose
Step 8	<pre>route-policy {in   out}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip-if)# route-policy out</p>	(Optional) Applies a routing policy to updates advertised to or received from a RIP neighbor.
Step 9	<pre>end or commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip-if)# end OR RP/0/RP0/CPU0:router(config-rip-if)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Creating a Route Policy for RIP

This task defines a route policy and shows how to attach it to an instance of a RIP process. Route policies can be used to:

- Control routes sent and received
- Control which routes are redistributed
- Control origination of the default route

A route policy definition consists of the **route-policy** command and *name* argument followed by a sequence of optional policy statements, and then closes with the **end-policy** command.

A route policy is not useful until it is applied to routes of a routing protocol.

### SUMMARY STEPS

- configure**
- route-policy** *name*
- set rip-metric** *number*
- end-policy**

5. **end**  
or  
**commit**
6. **configure**
7. **router rip**
8. **route-policy** *route-policy-name* {**in** | **out**}
9. **end**  
or  
**commit**]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy IN-IPv4	Defines a route policy and enters route-policy configuration mode.
Step 3	<b>set rip-metric</b> <i>number</i> RP/0/RP0/CPU0:router(config-rpl)# set rip metric 42	(Optional) Sets the RIP metric attribute.
Step 4	<b>end-policy</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end-policy	Ends the definition of a route policy and exits route-policy configuration mode.

	Command or Action	Purpose
Step 5	<pre>end or commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-rpl)# end  OR  RP/0/RP0/CPU0:router(config-rpl)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
Step 6	<pre>configure</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router# configure </p>	<p>Enters global configuration mode.</p>
Step 7	<pre>router rip</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config)# router rip </p>	<p>Configures a RIP routing process.</p>



	Command or Action	Purpose
Step 8	<pre>route-policy route-policy-name {in   out}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# route-policy IN in</p>	Applies a routing policy to updates advertised to or received from an RIP neighbor.
Step 9	<pre>end</pre> <p>or</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-rip)# end OR RP/0/RP0/CPU0:router(config-rip)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuration Examples for Implementing RIP on Cisco IOS XR Software

This section provides the following configuration examples:

- [Configuring a Basic RIP Configuration: Example, page RC-347](#)
- [Configuring RIP on the Provider Edge: Example, page RC-348](#)
- [Adjusting RIP Timers for each VRF Instance: Example, page RC-348](#)
- [Configuring Redistribution for RIP: Example, page RC-349](#)
- [Configuring Route Policies for RIP: Example, page RC-349](#)
- [Configuring Passive Interfaces and Explicit Neighbors for RIP: Example, page RC-350](#)
- [Controlling RIP Routes: Example, page RC-350](#)

### Configuring a Basic RIP Configuration: Example

The following example shows two Gigabit Ethernet interfaces configured with RIP.

```
interface GigabitEthernet0/6/0/0
  ipv4 address 172.16.0.1 255.255.255.0
!
```

```

interface GigabitEthernet0/6/0/2
  ipv4 address 172.16.2.12 255.255.255.0
!

router rip
  interface GigabitEthernet0/6/0/0
  !
  interface GigabitEthernet0/6/0/2
  !
!
```

## Configuring RIP on the Provider Edge: Example

The following example shows how to configure basic RIP on the PE with two VPN routing and forwarding (VRF) instances.

```

router rip
  interface GigabitEthernet0/6/0/0
  !
  vrf vpn0
    interface GigabitEthernet0/6/0/2
    !
  !
  vrf vpn1
    interface GigabitEthernet0/6/0/3
    !
  !
!
```

## Adjusting RIP Timers for each VRF Instance: Example

The following example shows how to adjust RIP timers for each VPN routing and forwarding (VRF) instance.

For VRF instance `vpn0`, the **timers basic** command sets updates to be broadcast every 10 seconds. If a router is not heard from in 30 seconds, the route is declared unusable. Further information is suppressed for an additional 30 seconds. At the end of the flush period (45 seconds), the route is flushed from the routing table.

For VRF instance `vpn1`, timers are adjusted differently: 20, 60, 60, and 70 seconds.

The **output-delay** command changes the interpacket delay for RIP updates to 10 milliseconds on `vpn1`. The default is that interpacket delay is turned off.

```

router rip
  interface GigabitEthernet0/6/0/0
  !
  vrf vpn0
    interface GigabitEthernet0/6/0/2
    !
    timers basic 10 30 30 35
  !
  vrf vpn1
    interface GigabitEthernet0/6/0/3
    !
    timers basic 20 60 60 70
    output-delay 10
  !
!
```

## Configuring Redistribution for RIP: Example

The following example shows how to redistribute Border Gateway Protocol (BGP) and static routes into RIP.

The RIP metric used for redistributed routes is determined by the route policy. If a route policy is not configured or the route policy does not set RIP metric, the metric is determined based on the redistributed protocol. For VPNv4 routes redistributed by BGP, the RIP metric set at the remote PE router is used, if valid.

In all other cases (BGP, IS-IS, OSPF, EIGRP, connected, static), the metric set by the **default-metric** command is used. If a valid metric cannot be determined, then redistribution does not happen.

```
route-policy ripred
  set rip-metric 5
end-policy
!

router rip
  vrf vpn0
    interface GigabitEthernet0/6/0/2
      !
      redistribute connected
      default-metric 3
    !
  vrf vpn1
    interface GigabitEthernet0/6/0/3
      !
      redistribute bgp 100 route-policy ripred
      redistribute static
      default-metric 3
    !
  !
!
```

## Configuring Route Policies for RIP: Example

The following example shows how to configure inbound and outbound route policies that are used to control which route updates are received by a RIP interface or sent out from a RIP interface.

```
prefix-set pf1
  10.1.0.0/24
end-set
!

prefix-set pf2
  150.10.1.0/24
end-set
!

route-policy policy_in
  if destination in pf1 then
    pass
  endif
end-policy
!

route-policy pass-all
  pass
end-policy
!
```

```

route-policy infil
  if destination in pf2 then
    add rip-metric 2
    pass
  endif
end-policy
!

router rip
interface GigabitEthernet0/6/0/0
  route-policy policy_in in
!
interface GigabitEthernet0/6/0/2
!
route-policy infil in
route-policy pass-all out

```

## Configuring Passive Interfaces and Explicit Neighbors for RIP: Example

The following example shows how to configure passive interfaces and explicit neighbors. When an interface is passive, it only accepts routing updates. In other words, no updates are sent out of an interface except to neighbors configured explicitly.

```

router rip
interface GigabitEthernet0/6/0/0
  passive-interface
!
interface GigabitEthernet0/6/0/2
!
neighbor 172.17.0.1
neighbor 172.18.0.5
!

```

## Controlling RIP Routes: Example

The following example shows how to use the **distance** command to install RIP routes in the Routing Information Base (RIB). The **maximum-paths** command controls the number of maximum paths allowed per RIP route.

```

router rip
interface GigabitEthernet0/6/0/0
  route-policy polin in
!
distance 110
maximum-paths 8
!

```

## Additional References

The following sections provide references related to implementing RIP on Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
RIP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference</i>
MPLS VPN support for RIP feature information	<i>Implementing MPLS Traffic Engineering on Cisco IOS XR Software module in Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>
Site of Origin (SoO) support for RIP feature information	<i>Implementing MPLS Traffic Engineering on Cisco IOS XR Software module in Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i>
Cisco CRS-1 router getting started material	<i>Cisco CRS-1 Carrier Routing System Getting Started Guide</i>
Information about user groups and task IDs	<i>Configuring AAA Services on Cisco IOS-XR Software module of Cisco IOS XR System Security Configuration Guide</i>

## Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
RFC 2453	RIP Version 2

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>



# Implementing Routing Policy on Cisco IOS XR Software

A routing policy instructs the router to inspect routes, filter them, and potentially modify their attributes as they are accepted from a peer, advertised to a peer, or redistributed from one routing protocol to another. Routing protocols make decisions to advertise, aggregate, discard, distribute, export, hold, import, redistribute and otherwise modify routes based on configured routing policy.

The routing policy language (RPL) has been designed to provide a single, straightforward language in which all routing policy needs can be expressed. RPL was designed to support large-scale routing configurations. It greatly reduces the redundancy inherent in previous routing policy configuration methods. RPL has been designed to streamline routing policy configuration, to reduce system resources required to store and process these configurations, and to simplify troubleshooting.



## Note

For more information about routing policy on the Cisco IOS XR software and complete descriptions of the routing policy commands listed in this module, see the “[Related Documents](#)” section of this module. To locate documentation for other commands that might appear during execution of a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing Routing Policy on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router.
Release 3.3.0	Support was added for EIGRP, RIP policy, clear-policy, debug, OSPF area-in, and OSPF area-out attach points.
Release 3.4.0	Support was added for the BGP next-hop attach point. Support was also added for null sets and global parameterization.
Release 3.5.0	The following features were added: <ul style="list-style-type: none"><li>• GNU Nano text editor for editing RPL policies.</li><li>• Enhanced prefix match functionality.</li><li>• Parameterization at attach points</li><li>• New ‘done’ disposition policy statement type</li></ul>
Release 3.6.0	No modification.

# Contents

- [Prerequisites for Implementing Routing Policy on Cisco IOS XR Software, page RC-354](#)
- [Information About Implementing Routing Policy on Cisco IOS XR Software, page RC-354](#)
- [How to Implement Routing Policy on Cisco IOS XR Software, page RC-408](#)
- [Configuration Examples for Implementing Routing Policy on Cisco IOS XR Software, page RC-413](#)
- [Additional References, page RC-416](#)

## Prerequisites for Implementing Routing Policy on Cisco IOS XR Software

The following are prerequisites for implementing Routing Policy on Cisco IOS XR Software:

- To use this command, you must be in a user group associated with a task group that includes the proper task IDs. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of the *Cisco IOS XR System Security Configuration Guide*.
- Border Gateway Protocol (BGP), integrated Intermediate System-to-Intermediate System (IS-IS), or Open Shortest Path First (OSPF) must be configured in your network.

## Information About Implementing Routing Policy on Cisco IOS XR Software

To implement RPL, you need to understand the following concepts:

- [Routing Policy Language, page RC-354](#)
- [Routing Policy Configuration Basics, page RC-363](#)
- [Policy Definitions, page RC-363](#)
- [Parameterization, page RC-364](#)
- [Semantics of Policy Application, page RC-365](#)
- [Policy Statements, page RC-370](#)
- [Attach Points, page RC-374](#)
- [Attached Policy Modification, page RC-406](#)
- [Nonattached Policy Modification, page RC-406](#)

## Routing Policy Language

This section contains the following information:

- [Routing Policy Language Overview, page RC-355](#)
- [Routing Policy Language Structure, page RC-355](#)



- [Routing Policy Language Components, page RC-360](#)
- [Routing Policy Language Usage, page RC-361](#)

## Routing Policy Language Overview

RPL was developed to support large-scale routing configurations. RPL has several fundamental capabilities that differ from those present in configurations oriented to traditional route maps, access lists, and prefix lists. The first of these capabilities is the ability to build policies in a modular form. Common blocks of policy can be defined and maintained independently. These common blocks of policy can then be applied from other blocks of policy to build complete policies. This capability reduces the amount of configuration information that needs to be maintained. In addition, these common blocks of policy can be parameterized. This parameterization allows for policies that share the same structure but differ in the specific values that are set or matched against to be maintained as independent blocks of policy. For example, three policies that are identical in every way except for the local preference value they set can be represented as one common parameterized policy that takes the varying local preference value as a parameter to the policy.

The policy language introduces the notion of sets. Sets are containers of similar data that can be used in route attribute matching and setting operations. Four set types exist: prefix-sets, community-sets, as-path-sets, and extcommunity-sets. These sets hold groupings of IPv4 or IPv6 prefixes, community values, AS path regular expressions, and extended community values, respectively. Sets are simply containers of data. Most sets also have an inline variant. An inline set allows for small enumerations of values to be used directly in a policy rather than having to refer to a named set. Prefix lists, community lists, and AS path lists must be maintained even when only one or two items are in the list. An inline set in RPL allows the user to place small sets of values directly in the policy body without having to refer to a named set.

Decision making, such as accept and deny, is explicitly controlled by the policy definitions themselves. RPL combines matching operators, which may use set data, with the traditional Boolean logic operators *and*, *or*, and *not* into complex conditional expressions. All matching operations return a true or false result. The execution of these conditional expressions and their associated actions can then be controlled by using simple *if then*, *elseif*, and *else* structures, which allow the evaluation paths through the policy to be fully specified by the user.

## Routing Policy Language Structure

This section describes the basic structure of RPL.

### Names

The policy language provides two kinds of persistent, namable objects: sets and policies. Definition of these objects is bracketed by beginning and ending command lines. For example, to define a policy named `test`, the configuration syntax would look similar to the following:

```
route-policy test
  [ . . . policy statements . . . ]
end-policy
```

Legal names for policy objects can be any sequence of the upper- and lowercase alphabetic characters; the numerals 0 to 9; and the punctuation characters period, hyphen, and underscore. A name must begin with a letter or numeral.

## Sets

In this context, the term set is used in its mathematical sense to mean an unordered collection of unique elements. The policy language provides sets as a container for groups of values for matching purposes. Sets are used in conditional expressions. The elements of the set are separated by commas. Null (empty) sets are allowed.

In the following example:

```
prefix-set backup-routes
  # currently no backup routes are defined
end-set
```

a condition such as:

```
if destination in backup-routes then
```

evaluates as FALSE for every route, because there is no match-condition in the prefix set that it satisfies.

Five kinds of sets exist: [as-path-set](#), [community-set](#), [extcommunity-set](#), [prefix-set](#), and [rd-set](#). You may want to perform comparisons against a small number of elements, such as two or three community values, for example. To allow for these comparisons, the user can enumerate these values directly. These enumerations are referred to as *inline sets*. Functionally, inline sets are equivalent to named sets, but allow for simple tests to be inline. Thus, comparisons do not require that a separate named set be maintained when only one or two elements are being compared. See the set types described in the following sections for the syntax. In general, the syntax for an inline set is a comma-separated list surrounded by parentheses as follows: (element-entry,element-entry,element-entry, ...element-entry), where element-entry is an entry of an item appropriate to the type of usage such as a prefix or a community value.

The following is an example using an inline community set:

```
route-policy sample-inline
  if community matches-any ([10..15]:100) then
    set local-preference 100
  endif
end-policy
```

The following is an equivalent example using the named set test-communities:

```
community-set test-communities
  10:100,
  11:100,
  12:100,
  13:100,
  14:100,
  15:100
end-set

route-policy sample
  if community matches-any test-communities then
    set local-preference 100
  endif
end-policy
```

Both of these policies are functionally equivalent, but the inline form does not require the configuration of the community set just to store the six values. You can choose the form appropriate to the configuration context. In the following sections, examples of both the named set version and the inline form are provided where appropriate.

## as-path-set

An AS path set comprises operations for matching an AS path attribute. The only matching operation is a regular expression match.

### Named Set Form

The named set form uses the **ios-regex** keyword to indicate the type of regular expression and requires single quotation marks around the regular expression.

The following is a sample definition of a named AS path set:

```
as-path-set aset1
  ios-regex '_42$',
  ios-regex '_127$'
end-set
```

This AS path set comprises two elements. When used in a matching operation, this AS path set matches any route whose AS path ends with either the autonomous system (AS) number 42 or 127.

To remove the named AS path set, use the **no as-path-set aset1** command-line interface (CLI) command.

### Inline Set Form

The inline set form is a parenthesized list of comma-separated expressions, as follows:

```
(ios-regex '_42$', ios-regex '_127$')
```

This set matches the same AS paths as the previously named set, but does not require the extra effort of creating a named set separate from the policy that uses it.

## community-set

A community-set holds community values for matching against the BGP community attribute. A community is a 32-bit quantity. Integer community values *must* be split in half and expressed as two unsigned decimal integers in the range from 0 to 65535, separated by a colon. Single 32-bit community values are not allowed. The following is the named set form:

### Named Set Form

```
community-set cset1
  12:34,
  12:56,
  12:78,
  internet
end-set
```

### Inline Set Form

```
(12:34, 12:56, 12:78)
($as:34, $as:$tag1, 12:78, internet)
```

The inline form of a community-set also supports parameterization. Each 16-bit portion of the community may be parameterized. See the [“Parameterization” section on page RC-364](#) for more information.

RPL provides symbolic names for the standard well-known community values: internet is 0:0, no-export is 65535:65281, no-advertise is 65535:65282, and local-as is 65535:65283.

RPL also provides a facility for using *wildcards* in community specifications. A wildcard is specified by inserting an asterisk (\*) in place of one of the 16-bit portions of the community specification; the wildcard indicates that any value for that portion of the community matches. Thus, the following policy matches all communities in which the autonomous system part of the community is 123:

```
community-set cset3
  123:*
end-set
```

Every community set must contain at least one community value. Empty community sets are invalid and are rejected.

## extcommunity-set

An extended community-set is analogous to a community-set except that it contains extended community values instead of regular community values. It also supports named forms and inline forms. There are three types of extended community sets: cost, soo, and rt.

As with community sets, the inline form supports parameterization within parameterized policies. Either portion of the extended community value can be parameterized.



### Note

*Wildcards* are not supported in extended communities.

Every extended community-set must contain at least one extended community value. Empty extended community-sets are invalid and rejected.

The following are syntactic examples:

### Named Form for Extcommunity-set Cost

A cost set is an extcommunity set used to store cost extended community type communities:

```
extcommunity-set cost a_cost_set
  IGP:1:10
end-set
```

### Named Form for Extcommunity-set RT

An rt set is an extcommunity set used to store route target extended community type communities:

```
extcommunity-set rt a_rt_set
  1.2.3.4:666
  1234:666,
  1.2.3.4:777,
  4567:777
end-set
```

Inline Set Form for Extcommunity-set RT

```
(1.2.3.4:666, 1234:666, 1.2.3.4:777, 4567:777)
($ipaddr:666, 1234:$tag, 1.2.3.4:777, $tag2:777)
```

### Named Form for Extcommunity-set Soo

A soo set is an extcommunity set used to store Site-of-Origin (SoO) extended community type communities:

```
extcommunity-set soo a_soo_set
  1.1.1.1:100,
  100:200
end-set
```

## prefix-set

A prefix-set holds IPv4 or IPv6 prefix match specifications, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. The address is a standard dotted-decimal IPv4 or colon-separated hexadecimal IPv6 address. The mask length, if present, is a nonnegative decimal integer in the range from 0 to 32 (0 to 128 for IPv6) following the address and separated from it by a slash. The optional minimum matching length follows the address and optional mask length and is expressed as the keyword **ge** (mnemonic for **greater than or equal to**), followed by a nonnegative decimal integer in the range from 0 to 32 (0 to 128 for IPv6). The optional maximum matching length follows the rest and is expressed by the keyword **le** (mnemonic for **less than or equal to**), followed by yet another nonnegative decimal integer in the range from 0 to 32 (0 to 128 for IPv6). A syntactic shortcut for specifying an exact length for prefixes to match is the **eq** keyword (mnemonic for **equal to**).

If a prefix match specification has no mask length, then the default mask length is 32 for IPv4 and 128 for IPv6. The default minimum matching length is the mask length. If a minimum matching length is specified, then the default maximum matching length is 32 for IPv4 and 128 for IPv6. Otherwise, if neither minimum nor maximum is specified, the default maximum is the mask length.

The prefix-set itself is a comma-separated list of prefix match specifications. The following are examples:

```
prefix-set legal-ipv4-prefix-examples
  10.0.1.1,
  10.0.2.0/24,
  10.0.3.0/24 ge 28,
  10.0.4.0/24 le 28,
  10.0.5.0/24 ge 26 le 30,
  10.0.6.0/24 eq 28,
  10.0.7.2/32 ge 16 le 24,
  10.0.8.0/26 ge 8 le 16
end-set
```

```
prefix-set legal-ipv6-prefix-examples
  2001:0:0:1::/64,
  2001:0:0:2::/64 ge 96,
  2001:0:0:2::/64 ge 96 le 100,
  2001:0:0:2::/64 eq 100
end-set
```

The first element of the prefix-set matches only one possible value, 10.0.1.1/32 or the host address 10.0.1.1. The second element matches only one possible value, 10.0.2.0/24. The third element matches a range of prefix values, from 10.0.3.0/28 to 10.0.3.255/32. The fourth element matches a range of values, from 10.0.4.0/24 to 10.0.4.240/28. The fifth element matches prefixes in the range from 10.0.5.0/26 to 10.0.5.252/30. The sixth element matches any prefix of length 28 in the range from 10.0.6.0/28 through 10.0.6.240/28. The seventh element matches any prefix of length 32 in the range 10.0.[0..255].2/32 (from 10.0.0.2/32 to 10.0.255.2). The eighth element matches any prefix of length 26 in the range 10.[0..255].8.0/26 (from 10.0.8.0/26 to 10.255.8.0/26).

The following prefix-set consists entirely of invalid prefix match specifications:

```
prefix-set ILLEGAL-PREFIX-EXAMPLES
  10.1.1.1 ge 16,
  10.1.2.1 le 16,
  10.1.3.0/24 le 23,
  10.1.4.0/24 ge 33,
  10.1.5.0/25 ge 29 le 28
end-set
```

Neither the minimum length nor maximum length is valid without a mask length. For IPv4, the minimum length must be less than 32, the maximum length of an IPv4 prefix. For IPv6, the minimum length must be less than 128, the maximum length of an IPv6 prefix. The maximum length must be equal to or greater than the minimum length.

## rd-set

An rd-set is used to create a set with route distinguisher (RD) elements. An RD set is a 64-bit value prepended to an IPv4 address to create a globally unique Border Gateway Protocol (BGP) VPN IPv4 address.

You can define RD values with the following commands:

- *a.b.c.d:m:\**—BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2:255.255.0.0:\*
- *a.b.c.d/m:n*—BGP VPN RD in IPv4 format with a mask. For example, 10.0.0.2:255.255.0.0:666.
- *a.b.c.d:\** —BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2:255.255.0.0.
- *a.b.c.d:n*— BGP VPN RD in IPv4 format. For example, 10.0.0.2:666.
- *asn:\**— BGP VPN RD in ASN format with a wildcard character. For example, 10002:255.255.0.0.
- *asn:n*—BGP VPN RD in ASN format. For example, 10002:666.

The following is an example of an rd-set:

```
rd-set rdset1
  10.0.0.0/8:*,
  10.0.0.0/8:777,
  10.0.0.0:*,
  10.0.0.0:777,
  65000:*,
  65000:777
end-set
```

## Routing Policy Language Components

Four main components in the routing policy language are involved in defining, modifying, and using policies: the configuration front end, policy repository, execution engine, and policy clients themselves.

The configuration front end (CLI) is the mechanism to define and modify policies. This configuration is then stored on the router using the normal storage means and can be displayed using the normal configuration **show** commands.

The second component of the policy infrastructure, the policy repository, has several responsibilities. First, it compiles the user-entered configuration into a form that the execution engine can understand. Second, it performs much of the verification of policies; and it ensures that defined policies can actually be executed properly. Third, it tracks which attach points are using which policies so that when policies are modified the appropriate clients are properly updated with the new policies relevant to them.

The third component is the execution engine. This component is the piece that actually runs policies as the clients request. The process can be thought of as receiving a route from one of the policy clients and then executing the actual policy against the specific route data.

The fourth component is the policy clients (the routing protocols). This component calls the execution engine at the appropriate times to have a given policy be applied to a given route, and then perform some number of actions. These actions may include deleting the route if policy indicated that it should be dropped, passing along the route to the protocol decision tree as a candidate for the best route, or advertising a policy modified route to a neighbor or peer as appropriate.

## Routing Policy Language Usage

This section provides basic routing policy language usage examples. See the [“How to Implement Routing Policy on Cisco IOS XR Software” section on page RC-408](#) for detailed information on how to implement routing policy language.

### The *pass* policy

The following example shows how the policy accepts all presented routes without modifying the routes.

```
route-policy quickstart-pass
  pass
end-policy
```

### The *drop everything* policy

The following example shows how the policy explicitly rejects all routes presented to it. This type of policy is used to ignore everything coming from a specific peer.

```
route-policy quickstart-drop
  drop
end-policy
```

### Ignore routes with specific AS numbers in the path

The following example shows the policy definition in three parts. First, the **as-path-set** command defines three regular expressions to match against an AS path. Second, the **route-policy** command applies the AS path set to a route. If the AS path attribute of the route matches the regular expression defined with the **as-path-set** command, the protocol refuses the route. Third, the route policy is attached to BGP neighbor 10.0.1.2. BGP consults the policy named `ignore_path_as` on routes received (imported) from neighbor 10.0.1.2.

```
as-path-set ignore_path
  ios-regex '_11_',
  ios-regex '_22_',
  ios-regex '_33_'
end-set

route-policy ignore_path_as
  if as-path in ignore_path then
    drop
  else
    pass
  endif
end-policy

router bgp 2
  neighbor 10.0.1.2 address-family ipv4 unicast policy ignore_path_as in
```

### Set community based on MED

The following example shows how the policy tests the MED of a route and modifies the community attribute of the route based on the value of the MED. If the MED value is 127, the policy adds the community 123:456 to the route. If the MED value is 63, the policy adds the value 123:789 to the community attribute of the route. Otherwise, the policy removes the community 123:123 from the route. In any case, the policy instructs the protocol to accept the route.

```
route-policy quickstart-med
  if med eq 127 then
    set community (123:456) additive
  elseif med eq 63 then
    set community (123:789) additive
  else
    delete community in (123:123)
  endif
  pass
end-policy
```

### Set local preference based on community

The following example shows how the community-set named quickstart-communities defines community values. The route policy named quickstart-localpref tests a route for the presence of the communities specified in the quickstart-communities community set. If any of the community values are present in the route, the route policy sets the local preference attribute of the route to 31. In any case, the policy instructs the protocol to accept the route.

```
community-set quickstart-communities
  987:654,
  987:543,
  987:321,
  987:210
end-set

route-policy quickstart-localpref
  if community matches-any quickstart-communities then
    set local-preference 31
  endif
  pass
end-policy
```

### Persistent Remarks

The following example shows how comments are placed in the policy to clarify the meaning of the entries in the set and the statements in the policy. The remarks are persistent, meaning they remain attached to the policy. For example, remarks are displayed in the output of the **show running-config** command. Adding remarks to the policy makes the policy easier to understand, modify at a later date, and troubleshoot if an unexpected behavior occurs.

```
prefix-set rfc1918
  # These are the networks defined as private in RFC1918 (including
  # all subnets thereof)
  10.0.0.0/8 ge 8,
  172.16.0.0/12 ge 12,
  192.168.0.0/16 ge 16
end-set

route-policy quickstart-remarks
  # Handle routes to RFC1918 networks
  if destination in rfc1918 then
    # Set the community such that we do not export the route
    set community (no-export) additive
```



```
endif
end-policy
```

## Routing Policy Configuration Basics

Route policies comprise series of statements and expressions that are bracketed with the **route-policy** and **end-policy** keywords. Rather than a collection of individual commands (one for each line), the statements within a route policy have context relative to each other. Thus, instead of each line being an individual command, each policy or set is an independent configuration object that can be used, entered, and manipulated as a unit.

Each line of a policy configuration is a logical subunit. At least one new line must follow the **then**, **else**, and **end-policy** keywords. A new line must also follow the closing parenthesis of a parameter list and the name string in a reference to an AS path set, community set, extended community set, or prefix set. At least one new line must precede the definition of a route policy, AS path set, community set, extended community set, or prefix set. One or more new lines can follow an action statement. One or more new lines can follow a comma separator in a named AS path set, community set, extended community set, or prefix set. A new line must appear at the end of a logical unit of policy expression and may not appear anywhere else.

## Policy Definitions

Policy definitions create named sequences of policy statements. A policy definition consists of the CLI **route-policy** keyword followed by a name, a sequence of policy statements, and the **end-policy** keyword. For example, the following policy drops any route it encounters:

```
route-policy drop-everything
  drop
end-policy
```

The name serves as a handle for binding the policy to protocols. To remove a policy definition, issue the **no route-policy name** command.

Policies may also refer to other policies such that common blocks of policy can be reused. This reference to other policies is accomplished by using the **apply** statement, as shown in the following example:

```
route-policy check-as-1234
  if as-path passes-through '1234.5' then
    apply drop-everything
  else
    pass
  endif
end-policy
```

The **apply** statement indicates that the policy drop-everything should be executed if the route under consideration passed through autonomous system 1234.5 before it is received. If a route that has autonomous system 1234.5 in its AS path is received, the route is dropped; otherwise, the route is accepted without modification. This policy is an example of a hierarchical policy. Thus, the semantics of the **apply** statement are just as if the applied policy were cut and pasted into the applying policy:

```
route-policy check-as-1234-prime
  if as-path passes-through '1234.5' then
    drop
  else
    pass
  endif
end-policy
```

You may have as many levels of hierarchy as desired. However, many levels may be difficult to maintain and understand.

## Parameterization

In addition to supporting reuse of policies using the **apply** statement, policies can be defined that allow for parameterization of some of the attributes. The following example shows how to define a parameterized policy named `param-example`. In this case, the policy takes one parameter, `$mytag`. Parameters always begin with a dollar sign and consist otherwise of any alphanumeric characters. Parameters can be substituted into any attribute that takes a parameter.

In the following example, a 16-bit community tag is used as a parameter:

```
route-policy param-example ($mytag)
    set community (1234:$mytag) additive
end-policy
```

This parameterized policy can then be reused with different parameterization, as shown in the following example. In this manner, policies that share a common structure but use different values in some of their individual statements can be modularized. For details on which attributes can be parameterized, see the individual attribute sections.

```
route-policy origin-10
    if as-path originates-from '10.5' then
        apply param-example(10.5)
    else
        pass
    endif
end-policy

route-policy origin-20
    if as-path originates-from '20.5' then
        apply param-example(20.5)
    else
        pass
    endif
end-policy
```

The parameterized policy `param-example` provides a policy definition that is expanded with the values provided as the parameters in the `apply` statement. Note that the policy hierarchy is always maintained. Thus, if the definition of `param-example` changes, then the behavior of `origin_10` and `origin_20` changes to match.

The effect of the `origin-10` policy is that it adds the community `1234:10` to all routes that pass through this policy and have an AS path indicating the route originated from autonomous system 10. The `origin-20` policy is similar except that it adds to community `1234:20` for routes originating from autonomous system 20.

## Parameterization at Attach Points

In addition to supporting parameterization using the `apply` statement described in the [“Parameterization” section on page RC-364](#), policies can also be defined that allow for parameterization of some of the attributes at attach points. Parameterization at BGP neighbor inbound and neighbor outbound attach points is supported.

In the example below, we define a parameterized policy "param-example". In this example, the policy takes two parameters "\$mymed" and "\$prefixset". Parameters always begin with a dollar sign, and consist otherwise of any alphanumeric characters. Parameters can be substituted into any attribute that takes a parameter. In this example we are passing a MED value and prefix set name as parameters.

```
route-policy param-example ($mymed, $prefixset)
  if destination in $prefixset then
    set med $mymed
  endif
end-policy
```

This parameterized policy can then be reused with different parameterizations as shown in the example below. In this manner, policies that share a common structure but use different values in some of their individual statements can be modularized. For details on which attributes can be parameterized, see the individual attributes for each protocol.

```
router bgp 2
  neighbor 10.1.1.1
    remote-as 3
    address-family ipv4 unicast
      route-policy param-example(10, prefix_set1)
      route-policy param-example(20, prefix_set2)
```

The parameterized policy param-example provides a policy definition that is expanded with the values provided as the parameters in the neighbor route-policy in and out statement.

## Global Parameterization

RPL supports the definition of systemwide global parameters that can be used inside policy definition. Global parameters can be configured as follows:

```
Policy-global
  glbpathtype 'ebgp'
  glbtag '100'
end-global
```

The global parameter values can be used directly inside a policy definition similar to the local parameters of parameterized policy. In the following example, the *globalparam* argument, which makes use of the global parameters *gpathtype* and *glbtag*, is defined for a nonparameterized policy.

```
route-policy globalparam
  if path-type is $glbpathtype then
    set tag $glbtag
  endif
end-policy
```

When a parameterized policy has a parameter name "collision" with a global parameter name, parameters local to policy definition take precedence, effectively masking off global parameters. In addition, a validation mechanism is in place to prevent the deletion of a particular global parameter if it is referred by any policy.

## Semantics of Policy Application

This section discusses how routing policies are evaluated and applied. The following concepts are discussed:

- [Boolean Operator Precedence, page RC-366](#)
- [Multiple Modifications of the Same Attribute, page RC-366](#)

- [When Attributes Are Modified, page RC-367](#)
- [Default Drop Disposition, page RC-367](#)
- [Control Flow, page RC-368](#)
- [Policy Verification, page RC-368](#)

## Boolean Operator Precedence

Boolean expressions are evaluated in order of operator precedence, from left to right. The highest precedence operator is *not*, followed by *and*, and then *or*. The following expression:

```
med eq 10 and not destination in (10.1.3.0/24) or community matches-any ([10..25]:35)
```

if fully parenthesized to display the order of evaluation, would look like this:

```
(med eq 10 and (not destination in (10.1.3.0/24))) or community matches-any ([10..25]:35)
```

The inner *not* applies only to the destination test; the *and* combines the result of the *not* expression with the Multi Exit Discriminator (MED) test; and the *or* combines that result with the community test. If the order of operations are rearranged:

```
not med eq 10 and destination in (10.1.3.0/24) or community matches-any ([10..25]:35)
```

then the expression, fully parenthesized, would look like the following:

```
((not med eq 10) and destination in (10.1.3.0/24)) or community matches-any ([10..25]:35)
```

## Multiple Modifications of the Same Attribute

When a policy replaces the value of an attribute multiple times, the last assignment wins because all actions are executed. Because the MED attribute in BGP is one unique value, the last value to which it gets set to wins. Therefore, the following policy results in a route with a MED value of 12:

```
set med 9
set med 10
set med 11
set med 12
```

This example is trivial, but the feature is not. It is possible to write a policy that effectively changes the value for an attribute. For example:

```
set med 8
if community matches-any cs1 then
    set local-preference 122
    if community matches-any cs2 then
        set med 12
    endif
endif
```

The result is a route with a MED of 8, unless the community list of the route matches both cs1 and cs2, in which case the result is a route with a MED of 12.

In the case in which the attribute being modified can contain only one value, it is easy to think of this case as the last statement wins. However, a few attributes can contain multiple values and the result of multiple actions on the attribute is cumulative rather than as a replacement. The first of these cases is the use of the **additive** keyword on community and extended community evaluation. Consider a policy of the form:

```
route-policy community-add
    set community (10:23)
```

```
    set community (10:24) additive
    set community (10:25) additive
end-policy
```

This policy sets the community string on the route to contain all three community values: 10:23, 10:24, and 10:25.

The second of these cases is AS path prepending. Consider a policy of the form:

```
route-policy prepend-example
  prepend as-path 2.5 3
  prepend as-path 666.5 2
end-policy
```

This policy prepends 666.5 666.5 2.5 2.5 2.5 to the AS path. This prepending is a result of all actions being taken and to the AS path being an attribute that contains an array of values rather than a simple scalar value.

## When Attributes Are Modified

A policy does not modify route attribute values until all tests have been completed. In other words, comparison operators always run on the initial data in the route. Intermediate modifications of the route attributes do not have a cascading effect on the evaluation of the policy. Take the following example:

```
if med eq 12 then
  set med 42
  if med eq 42 then
    drop
  endif
endif
endif
```

This policy never executes the drop statement because the second test (med eq 42) sees the original, unmodified value of the MED in the route. Because the MED has to be 12 to get to the second test, the second test always returns false.

## Default Drop Disposition

All route policies have a default action to drop the route under evaluation unless the route has been modified by a policy action or explicitly passed. Applied (nested) policies implement this disposition as though the applied policy were pasted into the point where it is applied.

Consider a policy to allow all routes in the 10 network and set their local preference to 200 while dropping all other routes. You might write the policy as follows:

```
route-policy two
  if destination in (10.0.0.0/8 ge 8 le 32) then
    set local-preference 200
  endif
end-policy

route-policy one
  apply two
end-policy
```

It may appear that policy one drops all routes because it neither contains an explicit **pass** statement nor modifies a route attribute. However, the applied policy does set an attribute for some routes and this disposition is passed along to policy one. The result is that policy one passes routes with destinations in network 10, and drops all others.

## Control Flow

Policy statements are processed sequentially in the order in which they appear in the configuration. Policies that hierarchically reference other policy blocks are processed as if the referenced policy blocks had been directly substituted inline. For example, if the following policies are defined:

```
route-policy one
  set weight 100
end-policy

route-policy two
  set med 200
end-policy

route-policy three
  apply two
  set community (2:666) additive
end-policy

route-policy four
  apply one
  apply three
  pass
end-policy
```

Policy four could be rewritten in an equivalent way as follows:

```
route-policy four-equivalent
  set weight 100
  set med 200
  set community (2:666) additive
  pass
end-policy
```



### Note

---

The **pass** statement is not required and can be removed to represent the equivalent policy in another way.

---

## Policy Verification

Several different types of verification occur when policies are being defined and used.

### Range Checking

As policies are being defined, some simple verifications, such as range checking of values, is done. For example, the MED that is being set is checked to verify that it is in a proper range for the MED attribute. However, this range checking cannot cover parameter specifications because they may not have defined values yet. These parameter specifications are verified when a policy is attached to an attach point. The policy repository also verifies that there are no recursive definitions of policy, and that parameter numbers are correct. At attach time, all policies must be well formed. All sets and policies that they reference must be defined and have valid values. Likewise, any parameter values must also be in the proper ranges.

### Incomplete Policy and Set References

As long as a given policy is not attached at an attach point, the policy is allowed to refer to nonexistent sets and policies, which allows for freedom of workflow. You can build configurations that reference sets or policy blocks that are not yet defined, and then can later fill in those undefined policies and sets, thereby achieving much greater flexibility in policy definition. Every piece of policy you want to

reference while defining a policy need not exist in the configuration. Thus, a user can define a policy sample that references the policy bar using an **apply** statement even if the policy bar does not exist. Similarly, a user can enter a policy statement that refers to a nonexistent set.

However, the existence of all referenced policies and sets is enforced when a policy is attached. If you attempt to attach the policy sample with the reference to an undefined policy bar at an inbound BGP policy using the **neighbor 1.2.3.4 address-family ipv4 unicast policy sample in** command, the configuration attempt is rejected because the policy bar does not exist.

Likewise, you cannot remove a route policy or set that is currently in use at an attach point because this removal would result in an undefined reference. An attempt to remove a route policy or set that is currently in use results in an error message to the user.

A condition exists that is referred to as a null policy in which the policy bar exists but has no statements, actions, or dispositions in it. In other words, the policy bar does exist as follows:

```
route-policy bar
end-policy
```

This is a valid policy block. It effectively forces all routes to be dropped because it is a policy block that never modifies a route, nor does it include the pass statement. Thus, the default action of drop for the policy block is followed.

## Attached Policy Modification

Policies that are in use do, on occasion, need to be modified. Traditionally, configuration changes are done by completely removing the relevant configuration and then re-entering it. However, this allows for a window of time in which no policy is attached and the default action takes place. RPL provides a mechanism for an atomic change so that if a policy is redeclared, or edited using a text editor, the new configuration is applied immediately—which allows for policies that are in use to be changed without having a window of time in which no policy is applied at the given attach point.

## Verification of Attribute Comparisons and Actions

The policy repository knows which attributes, actions, and comparisons are valid at each attach point. When a policy is attached, these actions and comparisons are verified against the capabilities of that particular attach point. Take, for example, the following policy definition:

```
route-policy bad
  set med 100
  set level level-1-2
  set ospf-metric 200
end-policy
```

This policy attempts to perform actions to set the BGP attribute med, IS-IS attribute level, and OSPF attribute cost. The system allows you to define such a policy, but it does not allow you to attach such a policy. If you had defined the policy bad and then attempted to attach it as an inbound BGP policy using the BGP configuration statement **neighbor 1.2.3.4 address-family ipv4 unicast route-policy bad in** the system would reject this configuration attempt. This rejection results from the verification process checking the policy and realizing that while BGP could set the MED, it has no way of setting the level or cost as the level and cost are attributes of IS-IS and OSPF, respectively. Instead of silently omitting the actions that cannot be done, the system generates an error to the user. Likewise, a valid policy in use at an attach point cannot be modified in such a way as to introduce an attempt to modify a nonexistent attribute or to compare against a nonexistent attribute. The verifiers test for nonexistent attributes and reject such a configuration attempt.

## Policy Statements

Four types of policy statements exist: remark, disposition (drop and pass), action (set), and if (comparator).

### Remark

A remark is text attached to policy configuration but otherwise ignored by the policy language parser. Remarks are useful for documenting parts of a policy. The syntax for a remark is text that has each line prepended with a pound sign (#):

```
# This is a simple one-line remark.

# This
# is a remark
# comprising multiple
# lines.
```

In general, remarks are used between complete statements or elements of a set. Remarks are not supported in the middle of statements or within an inline set definition.

Unlike traditional !-comments in the CLI, RPL remarks persist through reboots and when configurations are saved to disk or a TFTP server and then loaded back onto the router.

### Disposition

If a policy modifies a route, by default the policy accepts the route. RPL provides a statement to force the opposite—the **drop** statement. If a policy matches a route and executes a drop, the policy does not accept the route. If a policy does not modify the route, by default the route is dropped. To prevent the route from being dropped, the **pass** statement is used.

The **drop** statement indicates that the action to take is to discard the route. When a route is dropped, no further execution of policy occurs. For example, if after executing the first two statements of a policy the **drop** statement is encountered, the policy stops and the route is discarded.



#### Note

---

All policies have a default **drop** action at the end of execution.

---

The **pass** statement allows a policy to continue executing even though the route has not been modified. When a policy has finished executing, any route that has been modified in the policy or any route that has received a pass disposition in the policy, successfully passes the policy and completes the execution. If route policy B\_rp is applied within route policy A\_rp, execution continues from policy A\_rp to policy B\_rp and back to policy A\_rp provided prefix is not dropped by policy B\_rp.

```
route-policy A_rp
  set community (10:10)
  apply B_rp
end-policy
!

route-policy B_rp
  if destination in (121.23.0.0/16 le 32, 155.12.0.0/16 le 32) then
    set community (121:155) additive
  endif
end-policy
!
```



By default, a route is **dropped** at the end of policy processing unless either the policy **modifies** a route attribute or it passes the route by means of an explicit **pass** statement. For example, if route-policy B is applied within route-policy A, then execution continues from policy A to policy B and back to policy A, provided the prefix is not dropped by policy B.

```
route-policy A
  if as-path neighbor-is '123' then
    apply B
    policy statement N
  end-policy
```

Whereas the following policies pass all routes that they evaluate.

```
route-policy PASS-ALL
  pass
end-policy
```

```
route-policy SET-LPREF
  set local-preference 200
end-policy
```

In addition to being implicitly dropped, a route may be dropped by an **explicit drop** statement. **Drop** statements cause a route to be dropped immediately so that no further policy processing is done. Note also that a **drop** statement overrides any previously processed **pass** statements or attribute modifications. For example, the following policy drops all routes. The first **pass** statement is executed, but is then immediately overridden by the **drop** statement. The second **pass** statement never gets executed.

```
route-policy DROP-EXAMPLE
  pass
  drop
  pass
end-policy
```

When one policy applies another, it is as if the applied policy were copied into the right place in the applying policy, and then the same drop-and-pass semantics are put into effect. For example, policies ONE and TWO are equivalent to policy ONE-PRIME:

```
route-policy ONE
  apply two
  if as-path neighbor-is '123' then
    pass
  endif
end-policy

route-policy TWO
  if destination in (10.0.0.0/16 le 32) then
    drop
  endif
end-policy

route-policy ONE-PRIME
  if destination in (10.0.0.0/16 le 32) then
    drop
  endif
  if as-path neighbor-is '123' then
    pass
  endif
end-policy
```

Because the effect of an **explicit drop** statement is immediate, routes in 10.0.0.0/16 le 32 are dropped without any further policy processing. Other routes are then considered to see if they were advertised by autonomous system 123. If they were advertised, they are passed; otherwise, they are implicitly dropped at the end of all policy processing.

The **done** statement indicates that the action to take is to stop executing the policy and accept the route. When encountering a **done** statement, the route is passed and no further policy statements are executed. All modifications made to the route prior to the **done** statement are still valid.

## Action

An action is a sequence of primitive operations that modify a route. Most actions, but not all, are distinguished by the **set** keyword. In a route policy, actions can be grouped together. For example, the following is a route policy comprising three actions:

```
route-policy actions
  set med 217
  set community (12:34) additive
  delete community in (12:56)
end-policy
```

## If

In its simplest form, an **if** statement uses a conditional expression to decide which actions or dispositions should be taken for the given route. For example:

```
if as-path in as-path-set-1 then
  drop
endif
```

The example indicates that any routes whose AS path is in the set as-path-set-1 are dropped. The contents of the **then** clause may be an arbitrary sequence of policy statements.

The following example contains two action statements:

```
if origin is igp then
  set med 42
  prepend as-path 73.5 5
endif
```

The CLI provides support for the **exit** command as an alternative to the **endif** command.

The **if** statement also permits an **else** clause, which is executed if the if condition is false:

```
if med eq 8 then
  set community (12:34) additive
else
  set community (12:56) additive
endif
```

The policy language also provides syntax, using the **elseif** keyword, to string together a sequence of tests:

```
if med eq 150 then
  set local-preference 10
elseif med eq 200 then
  set local-preference 60
elseif med eq 250 then
  set local-preference 110
else
  set local-preference 0
endif
```

The statements within an **if** statement may themselves be **if** statements, as shown in the following example:

```
if community matches-any (12:34,56:78) then
  if med eq 150 then
    drop
  endif
  set local-preference 100
endif
```

This policy example sets the value of the local preference attribute to 100 on any route that has a community value of 12:34 or 56:78 associated with it. However, if any of these routes has a MED value of 150, then these routes with either the community value of 12:34 or 56:78 and a MED of 150 are dropped.

## Boolean Conditions

In the previous section describing the **if** statement, all of the examples use simple Boolean conditions that evaluate to either true or false. RPL also provides a way to build compound conditions from simple conditions by means of Boolean operators.

Three Boolean operators exist: negation (**not**), conjunction (**and**), and disjunction (**or**). In the policy language, negation has the highest precedence, followed by conjunction, and then by disjunction. Parentheses may be used to group compound conditions to override precedence or to improve readability.

The following simple condition:

```
med eq 42
```

is true only if the value of the MED in the route is 42, otherwise it is false.

A simple condition may also be negated using the **not** operator:

```
not next-hop in (10.0.2.2)
```

Any Boolean condition enclosed in parentheses is itself a Boolean condition:

```
(destination in prefix-list-1)
```

A compound condition takes either of two forms. It can be a simple expression followed by the **and** operator, itself followed by a simple condition:

```
med eq 42 and next-hop in (10.0.2.2)
```

A compound condition may also be a simpler expression followed by the **or** operator and then another simple condition:

```
origin is igp or origin is incomplete
```

An entire compound condition may be enclosed in parentheses:

```
(med eq 42 and next-hop in (10.0.2.2))
```

The parentheses may serve to make the grouping of subconditions more readable, or they may force the evaluation of a subcondition as a unit.

In the following example, the highest-precedence **not** operator applies only to the destination test, the **and** operator combines the result of the **not** expression with the community test, and the **or** operator combines that result with the MED test.

```
med eq 10 or not destination in (10.1.3.0/24) and community matches-any
([12..34]:[56..78])
```

With a set of parentheses to express the precedence, the result is the following:

```
med eq 10 or ((not destination in (10.1.3.0/24)) and community matches-any  
([12..34]:[56..78]))
```

The following is another example of a complex expression:

```
(origin is igp or origin is incomplete or not med eq 42) and next-hop in (10.0.2.2)
```

The left conjunction is a compound condition enclosed in parentheses. The first simple condition of the inner compound condition tests the value of the origin attribute; if it is Interior Gateway Protocol (IGP), then the inner compound condition is true. Otherwise, the evaluation moves on to test the value of the origin attribute again, and if it is incomplete, then the inner compound condition is true. Otherwise, the evaluation moves to check the next component condition, which is a negation of a simple condition.

## apply

As discussed in the sections on policy definitions and parameterization of policies, the **apply** command executes another policy (either parameterized or unparameterized) from within another policy, which allows for the reuse of common blocks of policy. When combined with the ability to parameterize common blocks of policy, the **apply** command becomes a powerful tool for reducing repetitive configuration.

## Attach Points

Policies do not become useful until they are applied to routes, and for policies to be applied to routes they need to be made known to routing protocols. In BGP, for example, there are several situations where policies can be used, the most common of these is defining import and export policy. The policy attach point is the point in which an association is formed between a specific protocol entity, in this case a BGP neighbor, and a specific named policy. It is important to note that a verification step happens at this point. Each time a policy is attached, the given policy and any policies it may apply are checked to ensure that the policy can be validly used at that attach point. For example, if a user defines a policy that sets the IS-IS level attribute and then attempts to attach this policy as an inbound BGP policy, the attempt would be rejected because BGP routes do not carry IS-IS attributes. Likewise, when policies are modified that are in use, the attempt to modify the policy is verified against all current uses of the policy to ensure that the modification is compatible with the current uses.

Each protocol has a distinct definition of the set of attributes (commands) that compose a route. For example, BGP routes may have a community attribute, which is undefined in OSPF. Routes in IS-IS have a level attribute, which is unknown to BGP. Routes carried internally in the RIB may have a tag attribute.

When a policy is attached to a protocol, the protocol checks the policy to ensure the policy operates using route attributes known to the protocol. If the protocol uses unknown attributes, then the protocol rejects the attachment. For example, OSPF rejects attachment of a policy that tests the values of BGP communities.

The situation is made more complex by the fact that each protocol has access to at least two distinct route types. In addition to native protocol routes, for example BGP or IS-IS, some protocol policy attach points operate on RIB routes, which is the common central representation. Using BGP as an example, the protocol provides an attach point to apply policy to routes redistributed from the RIB to BGP. An attach point dealing with two different kinds of routes permits a mix of operations: RIB attribute operations for matching and BGP attribute operations for setting.

**Note**

The protocol configuration rejects attempts to attach policies that perform unsupported operations.

The following sections describe the protocol attach points, including information on the attributes (commands) and operations that are valid for each attach point.

- [BGP Policy Attach Points, page RC-375](#)
- [OSPF Policy Attach Points, page RC-395](#)
- [OSPFv3 Policy Attach Points, page RC-397](#)
- [IS-IS Policy Attach Points, page RC-398](#)
- [EIGRP Policy Attach Points, page RC-400](#)
- [RIP Policy Attach Points, page RC-403](#)

See *Cisco IOS XR Routing Command Reference* for more information on the attributes and operations.

## BGP Policy Attach Points

This section describes each of the BGP policy attach points and provides a summary of the BGP attributes and operators.

- [Aggregation, page RC-375](#)
- [Dampening, page RC-376](#)
- [Default Originate, page RC-377](#)
- [Neighbor Export, page RC-377](#)
- [Neighbor Import, page RC-378](#)
- [Network, page RC-378](#)
- [Redistribute, page RC-379](#)
- [Show BGP, page RC-379](#)
- [Table Policy, page RC-380](#)
- [Import, page RC-381](#)
- [Export, page RC-381](#)
- [Retain Route-Target, page RC-382](#)
- [Allocate-Label, page RC-383](#)
- [Neighbor-ORF, page RC-383](#)
- [Next-hop, page RC-384](#)
- [Clear-Policy, page RC-384](#)
- [Debug, page RC-384](#)

### Aggregation

The aggregation attach point generates an aggregate route to be advertised based on the conditional presence of subcomponents of that aggregate. Policies attached at this attach point are also able to set any of the valid BGP attributes on the aggregated routes. For example, the policy could set a community value or a MED on the aggregate that is generated. The specified aggregate is generated if any routes

evaluated by the named policy pass the policy. More specifics of the aggregate are filtered using the **suppress-route** keyword. Any actions taken to set attributes in the route affect attributes on the aggregate.

In the policy language, the configuration is controlled by which routes pass the policy. The suppress map was used to selectively filter or suppress specific components of the aggregate when the summary-only flag is not set. In other words, when the aggregate and more specific components are being sent, some of the more specific components can be filtered using a suppress map. In the policy language, this is controlled by selecting the route and setting the suppress flag. The attribute-map allowed the user to set specific attributes on the aggregated route. In the policy language, setting attributes on the aggregated route is controlled by normal action operations.

In the following example, the aggregate address 10.0.0.0/8 is generated if there are any component routes in the range 10.0.0.0/8 ge 8 le 25 except for 10.2.0.0/24. Because summary-only is not set, all components of the aggregate are advertised. However, the specific component 10.1.0.0 are suppressed.

```
route-policy sample
  if destination in (10.0.0.0/8 ge 8 le 25) then
    set community (10:33)
  endif
  if destination in (10.2.0.0/24) then
    drop
  endif
  if destination in (10.1.0.0/24) then
    suppress-route
  endif
end-policy

router bgp 2
address-family ipv4
  aggregate-address 10.0.0.0/8 route-policy sample
  .
  .
  .
```

The effect of aggregation policy on the attributes of the aggregate is cumulative. Every time an aggregation policy matches a more specific route, the set operations in the policy may modify the aggregate. The aggregate in the following example has a MED value that varies according to the number of more specific routes that comprise the aggregate.

```
route-policy bumping-aggregation
  set med +5
end-policy
```

If there are three matching more specific routes, the MED of the aggregate is the default plus 15; if there are seventeen more specific routes, the MED of the aggregate is the default plus 85.

The order that the aggregation policy is applied to prefix paths is deterministic but unspecified. That is, a given set of routes always appears in the same order, but there is no way to predict the order.

A drop in aggregation policy does not prevent generation of an aggregate, but it does prevent the current more specific route from contributing to the aggregate. If another more specific route gives the route a pass, the aggregate is generated. Only one more specific pass is required to generate an aggregate.

## Dampening

The dampening attach point controls the default route-dampening behavior within BGP. Unless overridden by a more specific policy on the associate peer, all routes in BGP apply the associated policy to set their dampening attributes.

The following policy sets dampening values for BGP IPv4 unicast routes. Those routes that are more specific than a /25 take longer to recover after they have been dampened than routes that are less specific than /25.

```
route-policy sample_damp
  if destination in (0.0.0.0/0 ge 25) then
    set dampening halflife 30 others default
  else
    set dampening halflife 20 others default
  endif
end-policy

router bgp 2
  address-family ipv4 unicast
    bgp dampening route-policy sample_damp
  .
  .
  .
```

## Default Originate

The default originate attach point allows the default route (0.0.0.0/0) to be conditionally generated and advertised to a peer, based on the presence of other routes. It accomplishes this configuration by evaluating the associated policy against routes in the Routing Information Base (RIB). If any routes pass the policy, the default route is generated and sent to the relevant peer.

The following policy generates and sends a default-route to the BGP neighbor 10.0.0.1 if any routes that match 10.0.0.0/8 ge 8 le 32 are present in the RIB.

```
route-policy sample-originate
  if rib-has-route in (10.0.0.0/8 ge 8 le 32) then
    pass
  endif
end-policy

router bgp 2
  neighbor 10.0.0.1
    remote-as 3
    address-family ipv4 unicast
    default-originate policy sample-originate
  .
  .
  .
```

## Neighbor Export

The neighbor export attach point selects the BGP routes to send to a given peer or group of peers. The routes are selected by running the set of possible BGP routes through the associated policy. Any routes that pass the policy are then sent as updates to the peer or group of peers. The routes that are sent may have had their BGP attributes altered by the policy that has been applied.

The following policy sends all BGP routes to neighbor 10.0.0.5. Routes that are tagged with any community in the range 2:100 to 2:200 are sent with a MED of 100 and a community of 2:666. The rest of the routes are sent with a MED of 200 and a community of 2:200.

```
route-policy sample-export
  if community matches-any (2:[100-200]) then
    set med 100
    set community (2:666)
  else
    set med 200
    set community (2:200)
```

```

endif
end-policy

router bgp 2
 neighbor 10.0.0.5
   remote-as 3
   address-family ipv4 unicast
     route-policy sample-export out
   .
   .
   .

```

## Neighbor Import

The neighbor import attach point controls the reception of routes from a specific peer. All routes that are received by a peer are run through the attached policy. Any routes that pass the attached policy are passed to the BGP Routing Information Base (BRIB) as possible candidates for selection as best path routes.

When a BGP import policy is modified, it is necessary to rerun all the routes that have been received from that peer against the new policy. The modified policy may now discard routes that were previously allowed through, allow through previously discarded routes, or change the way the routes are modified. A new configuration option in BGP (**bgp auto-policy-soft-reset**) that allows this modification to happen automatically in cases for which either soft reconfiguration is configured or the BGP route-refresh capability has been negotiated.

The following example shows how to receive routes from neighbor 10.0.0.1. Any routes received with the community 3:100 have their local preference set to 100 and their community tag set to 2:666. All other routes received from this peer have their local preference set to 200 and their community tag set to 2:200.

```

route-policy sample_import
 if community matches-any (3:100) then
   set local-preference 100
   set community (2:666)
 else
   set local-preference 200
   set community (2:200)
 endif
end-policy

router bgp 2
 neighbor 10.0.0.1
   remote-as 3
   address-family ipv4 unicast
     route-policy sample_import in
   .
   .
   .

```

## Network

The network attach point controls the injection of routes from the RIB into BGP. A route policy attached at this point is able to set any of the valid BGP attributes on the routes that are being injected.

The following example shows a route policy attached at the network attach point that sets the well-known community no-export for any routes more specific than /24:

```

route-policy NetworkControl
 if destination in (0.0.0.0/0 ge 25) then
   set community (no-export) additive
 endif
end-policy

```



```
router bgp 2
  address-family ipv4 unicast
    network 172.16.0.5/27 route-policy NetworkControl
```

## Redistribute

The redistribute attach point allows routes from other sources to be advertised by BGP. The policy attached at this point is able to set any of the valid BGP attributes on the routes that are being redistributed. Likewise, selection operators allow a user to control what route sources are being redistributed and which routes from those sources.

The following example shows how to redistribute all routes from OSPF instance 12 into BGP. If OSPF were carrying a default route, it is dropped. Routes carrying a tag of 10 have their local preference set to 300 and the community value of 2:666 and no-advertise attached. All other routes have their local preference set to 200 and a community value of 2:100 set.

```
route-policy sample_redistribute
  if destination in (0.0.0.0/0) then
    drop
  endif
  if tag eq 10 then
    set local-preference 300
    set community (2:666, no-advertise)
  else
    set local-preference 200
    set community (2:100)
  endif
end-policy

router bgp 2
  address-family ipv4 unicast
    redistribute ospf 12 route-policy sample_redistribute
  .
  .
```

## Show BGP

The show bgp attach point allows the user to display selected BGP routes that pass the given policy. Any routes that are not dropped by the attached policy are displayed in a manner similar to the output of the **show bgp** command.

In the following example, the **show bgp route-policy** command is used to display any BGP routes carrying a MED of 5:

```
route-policy sample-display
  if med eq 5 then
    pass
  endif
end-policy
!
show bgp route-policy sample-display
```

A **show bgp policy route-policy** command also exists, which runs all routes in the RIB past the named policy as if the RIB were an outbound BGP policy. This command then displays what each route looked like before it was modified and after it was modified, as shown in the following example:

```
RP/0/RP0/CPU0:router# show rpl route-policy test2

route-policy test2
  if (destination in (10.0.0.0/8 ge 8 le 32)) then
    set med 333
```

```

endif
end-policy
!
RP/0/RP0/CPU0:router# show bgp

BGP router identifier 10.0.0.1, local AS number 2
BGP main routing table version 11
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*> 10.0.0.0          10.0.1.2             10           0 3 ?
*> 10.0.0.0/9        10.0.1.2             10           0 3 ?
*> 10.0.0.0/10       10.0.1.2             10           0 3 ?
*> 10.0.0.0/11       10.0.1.2             10           0 3 ?
*> 10.1.0.0/16       10.0.1.2             10           0 3 ?
*> 10.3.30.0/24      10.0.1.2             10           0 3 ?
*> 10.3.30.128/25    10.0.1.2             10           0 3 ?
*> 10.128.0.0/9      10.0.1.2             10           0 3 ?
*> 10.255.0.0/24     10.0.101.2           1000         555      0 100 e
*> 10.255.64.0/24    10.0.101.2           1000         555      0 100 e
....

RP/0/RP0/CPU0:router# show bgp policy route-policy test2

10.0.0.0/8 is advertised to 10.0.101.2

Path info:
  neighbor:10.0.1.2      neighbor router id:10.0.1.2
  valid external best
Attributes after inbound policy was applied:
  next hop:10.0.1.2
  MET ORG AS
  origin:incomplete neighbor as:3 metric:10
  aspath:3
Attributes after outbound policy was applied:
  next hop:10.0.1.2
  MET ORG AS
  origin:incomplete neighbor as:3 metric:333
  aspath:2 3
...

```

## Table Policy

The table policy attach point allows the user to configure traffic-index values on routes as they are installed into the global routing table. This attach point supports the BGP policy accounting feature. BGP policy accounting uses the traffic indexes that are set on the BGP routes to track various counters. This way, router operators can select different sets of BGP route attributes using the matching operations and then set different traffic indexes for each different class of route they are interested in tracking.

The following example shows how to set the traffic index to 10 in IPv4 unicast routes that originated from autonomous system 10.33. Likewise, any IPv4 unicast routes that originated from autonomous system 11.60 have their traffic index set to 11 when they are installed into the FIB. These traffic indexes are then used to count traffic being forwarded on these routes inline cards by enabling the BGP policy accounting counters on the interfaces of interest.

```

route-policy sample-table
  if as-path originates-from '10.33' then
    set traffic-index 10
  elseif as-path originates-from '11.60' then
    set traffic-index 11

```

```

endif
end-policy

router bgp 2
  address-family ipv4 unicast
    table-policy sample-table
  .
  .
  .

```

## Import

The import attach point provides control over the import of routes from the global VPN IPv4 table to a particular VPN routing and forwarding (VRF) instance.

For Layer 3 VPN networks, provider edge (PE) routers learn of VPN IPv4 routes through the Multiprotocol Internal Border Gateway Protocol (MP-iBGP) from other PE routers and automatically filters out route announcements that do not contain route targets that match any import route targets of its VRFs.

This automatic route filtering happens without RPL configuration; however, to provide more control over the import of routes in a VRF, you can configure a VRF import policy.

The following example shows how to perform matches based on a route target extended community and then sets the next hop. If the route has route target value 10:91, then the next hop is set to 206.0.0.1. If the route has route target value 11:92, then the next hop is set to 206.0.0.2. If the route has Site-of-Origin (SoO) value 10:111111 or 10:111222, then the route is dropped. All other non-matching routes are dropped.

```

route-policy bgpvrf_import
  if extcommunity rt matches-any (10:91) then
    set next-hop 206.0.0.1
  elseif extcommunity rt matches-every (11:92) then
    set next-hop 206.0.0.2
  elseif extcommunity soo matches-any (10:111111, 10:111222) then
    pass
  endif
end-policy

vrf vrf_import
  address-family ipv4 unicast
    import route-policy bgpvrf_import
  .
  .
  .

```

## Export

The export attach point provides control over the export of routes from a particular VRF to a global VPN IPv4 table.

For Layer 3 VPN networks, export route targets are added to the VPN IPv4 routes when VRF IPv4 routes are converted into VPN IPv4 routes and advertised through the MP-iBGP to other PE routers (or flow from one VRF to another within a PE router).

A set of export route targets is configured with the VRF without RPL configuration; however, to set route targets conditionally, you can configure a VRF export policy.

The following example shows some match and set operations supported for the export route policy. If a route matches 206.92.1.0/24 then the route target extended community is set to 10:101, and the weight is set to 211. If the route does not match 206.92.1.0/24 but the origin of the route is egp, then the local

preference is set to 212 and the route target extended community is set to 10:101. If the route does not match those specified criteria, then the route target extended community 10:111222 is added to the route. In addition, RT 10:111222 is added to the route that matches any of the previous conditions as well.

```
route-policy bgpvrfl_export
  if destination in (206.92.1.0/24) then
    set extcommunity rt (10:101)
    set weight 211
  elseif origin is egp then
    set local-preference 212
    set extcommunity rt (10:101)
  endif
  set extcommunity rt (10:111222) additive
end-policy
```

```
vrf vrf-export
  address-family ipv4 unicast
    export route-policy bgpvrfl_export
  .
  .
  .
```

## Retain Route-Target

The retain route target attach point within BGP allows the specification of match criteria based only on route target extended community. The attach point is useful at the route reflector (RR) or at the Autonomous System Boundary Router (ASBR).

Typically, an RR has to retain all IPv4 VPN routes to peer with its PE routers. These PEs might require routers tagged with different route target IPv4 VPN routes resulting in non-scalable RRs. You can achieve scalability if you configure an RR to retain routes with a defined set of route target extended communities, and a specific set of VPNs to service.

Another reason to use this attach point is for an ASBR. ASBRs do not require that VRFs be configured, but need this configuration to retain the IPv4 VPN prefix information.

The following example shows how to configure the route policy retainer and apply it to the retain route target attach point. The route is accepted if the route contains route target extended communities 10:615, 10:6150, and 15.15.15.15.15:15. All other non-matching routes are dropped.

```
extcommunity-set rt rtset1
  0:615,
  10:6150,
  15.15.15.15.15:15
end-set

route-policy retainer
  if extcommunity rt matches-any rtset1 then
    pass
  endif
end-policy

router bgp 2
  address-family vpnv4 unicast
    retain route-target route-policy retainer
  .
  .
  .
```

## Allocate-Label

The allocate-label attach point provides increased control based on various attribute match operations. This attach point is typically used in inter-AS option C to decide whether the label should be allocated or not when sending updates to the neighbor for the IPv4 labeled unicast address family. The attribute setting actions supported are for pass and drop.

The following example shows how to configure a route policy that passes the prefix 0.0.0.0 with prefix length 0. Label allocation happens only if prefix 0.0.0.0 exists.

```
route-policy label_policy
  if destination in (0.0.0.0/0) then
    pass
  endif
end-policy

router bgp 2
  vrf vrf1
    rd auto
    address-family ipv4 unicast
      allocate-label route-policy label-policy
    .
    .
    .
```

## Neighbor-ORF

The neighbor-orf attach point provides the filtering of incoming BGP route updates using only prefix-based matching. In addition to using this as an inbound filter, the prefixes and disposition (drop or pass) are sent to upstream neighbors as an Outbound Route Filter (ORF) to allow them to perform filtering.

The following example shows how to configure a route policy orf-preset and apply it to the neighbor ORF attach point. The prefix of the route is dropped if it matches any prefix specified in orf-preset (211.105.1.0/24, 211.105.5.0/24, 211.105.11.0/24). In addition to this inbound filtering, BGP also sends these prefix entries to the upstream neighbor with a permit or deny so that the neighbor can filter updates before sending them on to their destination.

```
prefix-set orf-preset
  211.105.1.0/24,
  211.105.5.0/24,
  211.105.11.0/24
end-set

route-policy policy-orf
  if orf prefix in orf-preset then
    drop
  endif
  if orf prefix in (211.105.3.0/24, 211.105.7.0/24, 211.105.13.0/24) then
    pass
  endif

router bgp 2
  neighbor 1.1.1.1
    remote-as 3
    address-family ipv4 unicast
      orf route-policy policy-orf
    .
    .
    .
```

## Next-hop

The next-hop attach point provides increased control based on protocol and prefix-based match operations. The attach point is typically used to decide whether to act on a next-hop notification (up or down) event.

Support for next-hop tracking allows BGP to monitor reachability for routes in the Routing Information Base (RIB) that can directly affect BGP prefixes. The route policy at the BGP next-hop attach point helps limit notifications delivered to BGP for specific prefixes. The route policy is applied on RIB routes. Typically, route policies are used in conjunction with next-hop tracking to monitor non-BGP routes.

The following example shows how to configure BGP nexthop tracking feature using a route policy to monitor static or connected routes with the prefix 20.0.0.0 and prefix length 8.

```
route-policy nxthp_policy_A
  if destination in (20.0.0.0/8) and protocol in (static, connected) then
    pass
  endif
end-policy

router bgp 2
  address-family ipv4 unicast
    nexthop route-policy nxthp_policy_A
  .
  .
  .
```

## Clear-Policy

The clear-policy attach point provides increased control based on various AS path match operations when using a **clear bgp** command. This attach point is typically used to decide whether to clear BGP flap statistics based on AS-path-based match operations.

The following example shows how to configure a route policy where the in operator evaluates to true if one or more of the regular expression matches in the set my-as-set successfully match the AS path associated with the route. If it is a match, then the **clear** command clears the associated flap statistics.

```
as-path-set my-as-set
  ios-regex '_12$',
  ios-regex '_13$'
end-set

route-policy policy_a
  if as-path in my-as-set then
    pass
  else
    drop
  endif
end-policy

clear bgp ipv4 unicast flap-statistics route-policy policy_a
```

## Debug

The debug attach point provides increased control based on prefix-based match operations. This attach point is typically used to filter debug output for various BGP commands based on the prefix of the route.

The following example shows how to configure a route policy that will only pass the prefix 20.0.0.0 with prefix length 8; therefore, the debug output shows up only for that prefix.

```
route-policy policy_b
  if destination in (20.0.0.0/8) then
```

```

    pass
  else
    drop

  endif
end-policy

debug bgp update policy_b

```

## BGP Attributes and Operators

Table 1 summarizes the BGP attributes and operators.

**Table 1** *BGP Attributes and Operators*

Attach Point	Attribute	Match	Set
aggregation	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity cost	n/a	set set additive
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	set set + set -
	next-hop	in	set
	origin	is	set
	source	in	n/a
	suppress-route	n/a	suppress-route
	weight	n/a	set

**Table 1** BGP Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
allocate-label	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	n/a
	destination	in	n/a
	label	n/a	set
	local-preference	is, ge, le, eq	n/a
	med	is, eg, ge, le	n/a
	next-hop	in	n/a
	origin	is	n/a
	source	in	n/a
clear-policy	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a



**Table 1** *BGP Attributes and Operators (continued)*

<b>Attach Point</b>	<b>Attribute</b>	<b>Match</b>	<b>Set</b>
dampening	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	n/a
	dampening	n/a/	set dampening... to set values that control the dampening (see <a href="#">Dampening</a> , page RC-376)
	destination	in	n/a
	local-preference	is, ge, le, eq	n/a
	med	is, eg, ge, le	n/a
	next-hop	in	n/a
	origin	is	n/a
	source	in	n/a
debug	destination	in	n/a
default originate	med	n/a	set set + set -
	rib-has-route	rib-has-route	n/a

**Table 1** *BGP Attributes and Operators (continued)*

<b>Attach Point</b>	<b>Attribute</b>	<b>Match</b>	<b>Set</b>
export (VRF)	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity rt	is-empty matches-any matches-every	set additive delete-in delete-not-in delete-all
	extcommunity soo	is-empty matches-any matches-every	n/a
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	n/a
	next-hop	in	n/a
	origin	is	n/a
	source	in	n/a
	weight	n/a	set

**Table 1** *BGP Attributes and Operators (continued)*

<b>Attach Point</b>	<b>Attribute</b>	<b>Match</b>	<b>Set</b>
import (VRF)	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	n/a
	destination	in	n/a
	extcommunity rt	is-empty matches-any matches-every	n/a
	extcommunity soo	is-empty matches-any matches-every	n/a
	local-preference	is, ge, le, eq	n/a
	med	is, eg, ge, le	n/a
	next-hop	in	set set peer address set destination vrf
	origin	is	n/a
	source	in	n/a
neighbor-in	as-path	in is-local length neighbor-is originates-from passes-through unique-length	prepend prepend most-recent replace
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community community with 'peeras'	is-empty matches-any matches-every	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity cost	n/a	set set additive

Table 1 BGP Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
	extcommunity rt	is-empty matches-any matches-every	set additive delete-in delete-not-in delete-all
	extcommunity soo	is-empty matches-any matches-every	n/a
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	set set + set -
	next-hop	in	set set peer address
	origin	is	set
	source	in	n/a
	weight	n/a	set
neighbor-out	as-path	in is-local length neighbor-is originates-from passes-through unique-length	prepend prepend most-recent replace
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community community with 'peeras'	is-empty matches-any matches-every	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity cost	n/a	set set additive
	extcommunity rt	is-empty matches-any matches-every	set additive delete-in delete-not-in delete-all
	extcommunity soo	is-empty matches-any matches-every	n/a
	local-preference	is, ge, le, eq	set

**Table 1** *BGP Attributes and Operators (continued)*

Attach Point	Attribute	Match	Set
	med	is, eg, ge, le	set set + set - set max-unreachable <sup>1</sup> set igp-cost
	next-hop	in	set set self
	origin	is	set
	path-type	is	n/a
	rd	in	n/a
	source	in	n/a
	unsuppress-route	n/a	unsuppress-route
	vpn-distinguisher	n/a	set
neighbor-orf	orf-prefix	in	n/a
network	as-path	n/a	prepend
	community	n/a	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity cost	n/a	set set additive
	mpls-label <sup>2</sup>	route-has-label	n/a
	local-preference	n/a	set
	med	n/a	set set + set -
	next-hop	in	set
	origin	n/a	set
	route-type	is	
	tag	is, ge, le, eq	n/a
	weight	n/a	set
next-hop	destination	in	n/a
	protocol	is in	n/a
	source	in	n/a

**Table 1** BGP Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
redistribute	as-path	n/a	prepend
	community	n/a	set set additive delete in delete not in delete all
	destination	in	n/a
	extcommunity cost	n/a	set set additive
	local-preference	n/a	set
	med	n/a	set set + set -
	next-hop	in	set
	origin	n/a	set
	mpls-label <sup>2</sup>	route-has-label	n/a
	route-type	is	n/a
	tag	is, eq, ge, le	n/a
	weight	n/a	set
retain-rt	extcommunity rt	is-empty matches-any matches-every	n/a

**Table 1** *BGP Attributes and Operators (continued)*

<b>Attach Point</b>	<b>Attribute</b>	<b>Match</b>	<b>Set</b>
show	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	n/a
	destination	in	n/a
	extcommunity rt	is-empty matches-any matches-every	n/a
	extcommunity soo	is-empty matches-any matches-every	n/a
	med	is, eg, ge, le	n/a
	next-hop	in	n/a
	origin	is	n/a
source	in	n/a	

**Table 1** BGP Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
table-policy	as-path	in is-local length neighbor-is originates-from passes-through unique-length	n/a
	as-path-length	is, ge, le, eq	n/a
	as-path-unique-length	is, ge, le, eq	n/a
	community	is-empty matches-any matches-every	n/a
	local-preference	is, ge, le, eq	n/a
	destination	in	n/a
	med	is, eg, ge, le	n/a
	next-hop	in	n/a
	origin	is	n/a
	rib-metric	n/a	set
	source	in	n/a
	tag	n/a	set
traffic-index	n/a	set	

1. The max-reachable value is 4294967295.
2. The mpls-label attribute is implicit in the conditional syntax; it is implied by the *route-has-label* operator.

Some BGP route attributes are inaccessible from some BGP attach points for various reasons. For example, the **set med igp-cost only** command makes sense when there is a configured igp-cost to provide a source value. [Table 2](#) summarizes which operations are valid and where they are valid.

**Table 2** Restricted BGP Operations by Attach Point

	import	export	aggregation	redistribution
<b>prepend as-path</b>	eBGP only	eBGP only	n/a	n/a
<b>prepend as-path most-recent</b>	eBGP only	eBGP only	n/a	n/a
<b>replace as-path</b>	eBGP only	eBGP only	n/a	n/a
<b>set med igp-cost</b>	forbidden	eBGP only	forbidden	forbidden
<b>set weight</b>	n/a	forbidden	n/a	n/a
<b>suppress</b>	forbidden	forbidden	n/a	forbidden



## OSPF Policy Attach Points

This section describes each of the OSPF policy attach points and provides a summary of the OSPF attributes and operators.

- [Default-Information Originate, page RC-395](#)
- [Redistribute, page RC-395](#)
- [Area-in, page RC-396](#)
- [Area-out, page RC-396](#)

### Default-Information Originate

The default-information originate attach point allows the user to conditionally inject the default route 0.0.0.0/0 into the OSPF link-state database, which is done by evaluating the attached policy. If any routes in the local RIB pass the policy, then the default route is inserted into the link-state database.

The following example shows how to generate a default route if any of the routes that match 10.0.0.0/8 ge 8 le 25 are present in the RIB:

```
route-policy ospf-originate
  if rib-has-route in (10.0.0.0/8 ge 8 le 25) then
    pass
  endif
end-policy

router ospf 1
  default-information originate policy ospf-originate
  .
  .
  .
```

### Redistribute

The redistribute attach point within OSPF injects routes from other routing protocol sources into the OSPF link-state database, which is done by selecting the routes it wants to import from each protocol. It then sets the OSPF parameters of cost and metric type. The policy can control how the routes are injected into OSPF by using the **set metric-type** or **set ospf-metric** command.

The following example shows how to redistribute routes from IS-IS instance instance\_10 into OSPF instance 1 using the policy OSPF-redist. The policy sets the metric type to type-2 for all redistributed routes. IS-IS routes with a tag of 10 have their cost set to 100, and IS-IS routes with a tag of 20 have their OSPF cost set to 200. Any IS-IS routes not carrying a tag of either 10 or 20 are not be redistributed into the OSPF link-state database.

```
route-policy OSPF-redist
  set metric-type type-2
  if tag eq 10 then
    set ospf cost 100
  elseif tag eq 20 then
    set ospf cost 200
  else
    drop
  endif
end-policy
router ospf 1
  redistribute isis instance_10 policy OSPF-redist
  .
  .
  .
```

## Area-in

The area-in attach point within OSPF allows you to filter inbound OSPF type-3 summary link-state advertisements (LSAs). The attach point provides prefix-based matching and hence increased control for filtering type-3 summary LSAs.

The following example shows how to configure the prefix for OSPF summary LSAs. If the prefix matches any of 111.105.3.0/24, 111.105.7.0/24, 111.105.13.0/24, it is accepted. If the prefix matches any of 111.106.3.0/24, 111.106.7.0/24, 111.106.13.0/24, it is dropped.

```
route-policy OSPF-area-in
  if destination in (111.105.3.0/24, 111.105.7.0/24, 111.105.13.0/24) then
    drop
  endif
  if destination in (111.106.3.0/24, 111.106.7.0/24, 111.106.13.0/24) then
    pass
  endif
end-policy

router ospf 1
  area 1
    route-policy OSPF-area-in in
```

## Area-out

The area-out attach point within OSPF allows you to filter outbound OSPF type-3 summary LSAs. The attach point provides prefix-based matching and, hence, increased control for filtering type-3 summary LSAs.

The following example shows how to configure the prefix for OSPF summary LSAs. If the prefix matches any of 211.105.3.0/24, 211.105.7.0/24, 211.105.13.0/24, it is announced. If the prefix matches any of .105.3.0/24, 212.105.7.0/24, 212.105.13.0/24, it is dropped and not announced.

```
route-policy OSPF-area-out
  if destination in (211.105.3.0/24, 211.105.7.0/24, 211.105.13.0/24) then
    drop
  endif
  if destination in (212.105.3.0/24, 212.105.7.0/24, 212.105.13.0/24) then
    pass
  endif
end-policy

router ospf 1
  area 1
    route-policy OSPF-area-out out
```

## OSPF Attributes and Operators

[Table 3](#) summarizes the OSPF attributes and operators.

**Table 3** *OSPF Attributes and Operators*

Attach Point	Attribute	Match	Set
default-information originate	ospf-metric	n/a	set
	metric-type	n/a	set
	tag	n/a	set
	rib-has-route	in	n/a

**Table 3** OSPF Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
redistribute	destination	in	n/a
	metric-type	n/a	set
	ospf-metric	n/a	set
	next-hop	in	n/a
	mpls-label <sup>1</sup>	route-has-label	n/a
	route-type	is	n/a
	tag	is, eq, ge, le	set
area-in	destination	in	n/a
area-out	destination	in	n/a

1. The mpls-label attribute is implicit in the conditional syntax; it is implied by the *route-has-label* operator.

## OSPFv3 Policy Attach Points

This section describes each of the OSPFv3 policy attach points and provides a summary of the OSPFv3 attributes and operators.

- [Default-Information Originate, page RC-397](#)
- [Redistribute, page RC-397](#)

### Default-Information Originate

The default-information originate attach point allows the user to conditionally inject the default route 0::/0 into the OSPFv3 link-state database, which is done by evaluating the attached policy. If any routes in the local RIB pass the policy, then the default route is inserted into the link-state database.

The following example shows how to generate a default route if any of the routes that match 2001::/96 are present in the RIB:

```
route-policy ospfv3-originate
  if rib-has-route in (2001::/96) then
    pass
  endif
end-policy

router ospfv3 1
  default-information originate policy ospfv3-originate
  .
  .
```

### Redistribute

The redistribute attach point within OSPFv3 injects routes from other routing protocol sources into the OSPFv3 link-state database, which is done by selecting the route types it wants to import from each protocol. It then sets the OSPFv3 parameters of cost and metric type. The policy can control how the routes are injected into OSPFv3 by using the **metric type** command.

The following example shows how to redistribute routes from BGP instance 15 into OSPF instance 1 using the policy OSPFv3-redist. The policy sets the metric type to type-2 for all redistributed routes. BGP routes with a tag of 10 have their cost set to 100, and BGP routes with a tag of 20 have their OSPFv3 cost set to 200. Any BGP routes not carrying a tag of either 10 or 20 are not be redistributed into the OSPFv3 link-state database.

```
route-policy OSPFv3-redist
  set metric-type type-2
  if tag eq 10 then
    set extcommunity cost 100
  elseif tag eq 20 then
    set extcommunity cost 200
  else
    drop
  endif
end-policy

router ospfv3 1
  redistribute bgp 15 policy OSPFv3-redist
  .
  .
  .
```

## OSPFv3 Attributes and Operators

Table 4 summarizes the OSPFv3 attributes and operators.

**Table 4** OSPFv3 Attributes and Operators

Attach Point	Attribute	Match	Set
default-information originate	ospf-metric	n/a	set
	metric-type	n/a	set
	tag	n/a	set
	rib-has-route	in	n/a
redistribute	destination	in	n/a
	ospf-metric	n/a	set
	metric-type	n/a	set
	route-type	is	n/a
	tag	is, eq, ge, le	n/a

## IS-IS Policy Attach Points

This section describes each of the IS-IS policy attach points and provides a summary of the IS-IS attributes and operators.

- [Redistribute, page RC-399](#)
- [Default-Information Originate, page RC-399](#)

## Redistribute

The redistribute attach point within IS-IS allows routes from other protocols to be readvertised by IS-IS. The policy is a set of control structures for selecting the types of routes that a user wants to redistribute into IS-IS. The policy can also control which IS-IS level the routes are injected into and at what metric values.

The following example shows how to redistribute routes from IS-IS instance 1 into IS-IS instance instance\_10 using the policy ISIS-redist. This policy sets the level to level-1-2 for all redistributed routes. OSPF routes with a tag of 10 have their metric set to 100, and IS-IS routes with a tag of 20 have their IS-IS metric set to 200. Any IS-IS routes not carrying a tag of either 10 or 20 are not be redistributed into the IS-IS database.

```
route-policy ISIS-redist
  set level level-1-2
  if tag eq 10 then
    set metric 100
  elseif tag eq 20 then
    set metric 200
  else
    drop
  endif
end-policy

router isis instance_10
  address-family ipv4 unicast
    redistribute ospf 1 policy ISIS-redist
  .
  .
  .
```

## Default-Information Originate

The default-information originate attach point within IS-IS allows the default route 0.0.0.0/0 to be conditionally injected into the IS-IS route database.

The following example shows how to generate an IPv4 unicast default route if any of the routes that match 10.0.0.0/8 ge 8 le 25 is present in the RIB. The cost of the IS-IS route is set to 100 and the level is set to level-1-2 on the default route that is injected into the IS-IS database.

```
route-policy isis-originate
  if rib-has-route in (10.0.0.0/8 ge 8 le 25) then
    set metric 100
    set level level-1-2
  endif
end-policy

router isis instance_10
  address-family ipv4 unicast
    default-information originate policy isis_originate
  .
  .
  .
```

## Inter-area-propagate

The inter-area-propagate attach point within IS-IS allows the prefixes to be conditionally propagated from one level to another level within the same IS-IS instance.

The following example shows how to allow prefixes to be leaked from the level 1 LSP into the level 2 LSP if any of the prefixes match 10.0.0.0/8 ge 8 le 25.

```

route-policy isis-propagate
  if destination in (10.0.0.0/8 ge 8 le 25) then
    pass
  endif
end-policy

router isis instance_10
  address-family ipv4 unicast
    propagate level 1 into level 2 policy isis-propagate
  .
  .
  .

```

## IS-IS Attributes and Operators

Table 5 summarizes the IS-IS attributes and operators.

**Table 5** IS-IS Attributes and Operators

Attach Point	Attribute	Match	Set
redistribution	tag	is, eq, ge, le	set
	route-type	is	n/a
	destination	in	n/a
	next-hop	in	n/a
	mpls-label <sup>1</sup>	route-has-label	n/a
	level	n/a	set
	isis-metric	n/a	set
	metric-type	n/a	set
default-information originate	rib-has-route	in	n/a
	level	n/a	set
	isis-metric	n/a	set
	tag	n/a	set
inter-area-propagate	destination	in	n/a

1. The mpls-label attribute is implicit in the conditional syntax; it is implied by the *route-has-label* operator.

## EIGRP Policy Attach Points

This section describes each of the EIGRP policy attach points and provides a summary of the EIGRP attributes and operators.

- [Default-Accept-In, page RC-401](#)
- [Default-Accept-Out, page RC-401](#)
- [Policy-In, page RC-401](#)
- [Policy-Out, page RC-402](#)
- [If-Policy-In, page RC-402](#)
- [If-Policy-Out, page RC-402](#)
- [Redistribute, page RC-402](#)

## Default-Accept-In

The default-accept-in attach point allows you to set and reset the conditional default flag for EIGRP routes by evaluating the attached policy.

The following example shows a policy that sets the conditional default flag for all routes that match 10.0.0.0/8 and longer prefixes up to 10.0.0.0/25:

```
route-policy eigrp-cd-policy-in
  if destination in (10.0.0.0/8 ge 8 le 25) then
    pass
  endif
end-policy
!
router eigrp 100
  address-family ipv4
    default-information allowed in route-policy eigrp-cd-policy-in
  .
  .
  .
```

## Default-Accept-Out

The default-accept-out attach point allows you to set and reset the conditional default flag for EIGRP routes by evaluating the attached policy.

The following example shows a policy that sets the conditional default flag for all routes that match 100.10.0.0/16:

```
route-policy eigrp-cd-policy-out
  if destination in (200.10.0.0/16) then
    pass
  endif
end-policy
!
router eigrp 100
  address-family ipv4
    default-information allowed out route-policy eigrp-cd-policy-out
  .
  .
  .
```

## Policy-In

The policy-in attach point allows you to filter and modify inbound EIGRP routes. This policy is applied to all interfaces for which there is no interface inbound route policy.

The following example shows the command under EIGRP:

```
router eigrp 100
  address-family ipv4
    route-policy global-policy-in in
  .
  .
  .
```

## Policy-Out

The policy-out attach point allows you to filter and modify outbound EIGRP routes. This policy is applied to all interfaces for which there is no interface outbound route policy.

The following example shows the command under EIGRP:

```
router eigrp 100
  address-family ipv4
    route-policy global-policy-out out
  .
  .
  .
```

## If-Policy-In

The if-policy-in attach point allows you to filter routes received on a particular EIGRP interface. The following example shows an inbound policy for Packet-over-SONET/SDH (POS) interface 0/2/0/3:

```
router eigrp 100
  address-family ipv4
    interface POS0/2/0/3
      route-policy if-filter-policy-in in
    .
    .
    .
```

## If-Policy-Out

The if-policy-out attach point allows you to filter routes sent out on a particular EIGRP interface. The following example shows an outbound policy for Packet-over-SONET/SDH (POS) interface 0/2/0/3:

```
router eigrp 100
  address-family ipv4
    interface POS0/2/0/3
      route-policy if-filter-policy-out out
    .
    .
    .
```

## Redistribute

The redistribute attach point in EIGRP allows you to filter redistributed routes from other routing protocols and modify some routing parameters before installing the route in the EIGRP database. The following example shows a policy filter redistribution of RIP routes into EIGRP.

```
router-policy redistribute-rip
  if destination in (100.1.1.0/24) then
    set eigrp-metric 5000000 4000 150 30 2000
  else
    set tag 200
  endif
end-policy

router eigrp 100
  address-family ipv4
    redistribute rip route-policy redistribute-rip
  .
  .
  .
```



## EIGRP Attributes and Operators

Table 6 summarizes the EIGRP attributes and operators.

**Table 6** *EIGRP Attributes and Operators*

Attach Point	Attribute	Match	Set
default-accept-in	destination	in	n/a
default-accept-out	destination	in	n/a
if-policy-in	destination	in	n/a
	next-hop	in	n/a
	eigrp-metric	n/a	add, set
	tag	is, eq, ge, le	set
if-policy-out	destination	in	n/a
	next-hop	in	n/a
	protocol	is, in	n/a
	eigrp-metric	n/a	add, set
	tag	is, eq, ge, le	set
policy-in	destination	in	n/a
	next-hop	in	n/a
	eigrp-metric	n/a	add, set
	tag	is, eq, ge, le	set
policy-out	destination	in	n/a
	next-hop	in	n/a
	protocol	is, in	n/a
	eigrp-metric	n/a	add, set
	tag	is, eq, ge, le	set
redistribute	destination	in	n/a
	next-hop	in	n/a
	mpls-label <sup>1</sup>	route-has-label	n/a
	eigrp-metric	n/a	add, set
	route-type	is	n/a
	tag	is, eq, ge, le	set

1. The mpls-label attribute is implicit in the conditional syntax; it is implied by the *route-has-label* operator.

## RIP Policy Attach Points

This section describes each of the RIP policy attach points and provides a summary of the RIP attributes and operators.

- [Default-Information Originate, page RC-404](#)
- [Redistribute, page RC-404](#)

- [Global-Inbound, page RC-404](#)
- [Global-Outbound, page RC-404](#)
- [Interface-Inbound, page RC-405](#)
- [Interface-Outbound, page RC-405](#)

## Default-Information Originate

The default-information originate attach point allows you to conditionally inject the default route 0.0.0.0/0 into RIP updates by evaluating the attached policy. If any routes in the local RIB pass the policy, then the default route is inserted.

The following example shows how to generate a default route if any of the routes that match 10.0.0.0/8 ge 8 le 25 are present in the RIB:

```
route-policy rip-originate
  if rib-has-route in (10.0.0.0/8 ge 8 le 25) then
    pass
  endif
end-policy

router rip
  default-information originate route-policy rip-originate
```

## Redistribute

The redistribution attach point within RIP allows you to inject routes from other routing protocol sources into the RIP database.

The following example shows how to inject OSPF routes into RIP:

```
route-policy redistrib-ospf
  set rip-metric 5
end-policy

router rip
  redistribute ospf 1 route-policy redistrib-ospf
```

## Global-Inbound

The global-inbound attach point for RIP allows you to filter or update inbound RIP routes that match a route policy.

The following example shows how to filter the inbound RIP routes that match the route policy named rip-in:

```
router rip
  route-policy rip-in in
```

## Global-Outbound

The global-outbound attach point for RIP allows you to filter or update outbound RIP routes that match a route-policy.

The following example shows how to filter the outbound RIP routes that match the route policy named rip-out:

```
router rip
  route-policy rip-out out
```

## Interface-Inbound

The interface-inbound attach point allows you to filter or update inbound RIP routes that match a route policy for a specific interface.

The following example shows how to filter inbound RIP routes that match the route policy for interface 0/1/0/1:

```
router rip
  interface GigabitEthernet0/1/0/1
    route-policy rip-in in
```

## Interface-Outbound

The interface-outbound attach point allows you to filter or update outbound RIP routes that match a route policy for a specific interface.

The following example shows how to filter outbound RIP routes that match the route policy for interface 0/2/0/1:

```
router rip
  interface GigabitEthernet0/2/0/1
    route-policy rip-out out
```

## RIP Attributes and Operators

Table 7 summarizes the RIP attributes and operators.

**Table 7** *RIP Attributes and Operators*

Attach Point	Attribute	Match	Set
default-information originate	next-hop	n/a	set
	rip-metric	n/a	set
	rip-tag	n/a	set
	rib-has-route	in	n/a
global-inbound	destination	in	n/a
	next-hop	in	n/a
	rip-metric	n/a	add, set
global-outbound	destination	in	n/a
	protocol	is, in	n/a
	rip-metric	n/a	add, set
interface-inbound	destination	in	n/a
	next-hop	in	n/a
	rip-metric	n/a	add, set
interface-outbound	destination	in	n/a
	protocol	is, in	n/a
	rip-metric	n/a	add, set

Table 7 RIP Attributes and Operators (continued)

Attach Point	Attribute	Match	Set
redistribute	destination	in	n/a
	next-hop	in	set
	rip-metric	n/a	set
	rip-tag	n/a	set
	mpls-label <sup>1</sup>	route-has-label	n/a
	route-type	is	n/a
	tag	is, eq, ge, le	set

1. The mpls-label attribute is implicit in the conditional syntax; it is implied by the *route-has-label* operator.

## Attached Policy Modification

Policies that are in use do, on occasion, need to be modified. In the traditional configuration model, a policy modification would be done by completely removing the policy and re-entering it. However, this model allows for a window of time in which no policy is attached and default actions to be used, which is an opportunity for inconsistencies to exist. To close this window of opportunity, you can modify a policy in use at an attach point by respecifying it, which allows for policies that are in use to be changed, without having a window of time in which no policy is applied at the given attach point.



### Note

A route policy or set that is in use at an attach point cannot be removed because this removal would result in an undefined reference. An attempt to remove a route policy or set that is in use at an attach point results in an error message to the user.

## Nonattached Policy Modification

As long as a given policy is not attached at an attach point, the policy is allowed to refer to nonexistent sets and policies. Configurations can be built that reference sets or policy blocks that are not yet defined, and then later those undefined policies and sets can be filled in. This method of building configurations gives much greater flexibility in policy definition. Every piece of policy you want to reference while defining a policy need not exist in the configuration. Thus, you can define a policy sample1 that references a policy sample2 using an apply statement even if the policy sample2 does not exist. Similarly, you can enter a policy statement that refers to a nonexistent set.

However, the existence of all referenced policies and sets is enforced when a policy is attached. Thus, if a user attempts to attach the policy sample1 with the reference to an undefined policy sample2 at an inbound BGP policy using the statement **neighbor 1.2.3.4 address-family ipv4 unicast policy sample1 in**, the configuration attempt is rejected because the policy sample2 does not exist.

## Editing Routing Policy Configuration Elements

RPL is based on statements rather than on lines. That is, within the begin-end pair that brackets policy statements from the CLI, a new line is merely a separator, the same as a space character.

The CLI provides the means to enter and delete route policy statements. RPL provides a means to edit the contents of the policy between the begin-end brackets, using a text editor. The following text editors are available on Cisco IOS XR software for editing RPL policies:

- Nano (default)
- Emacs
- Vim

### Editing Routing Policy Configuration Elements Using the Nano Editor

To edit the contents of a routing policy using the Nano editor, use the following CLI command in EXEC mode:

```
edit route-policy name nano
```

A copy of the route policy is copied to a temporary file and the editor is launched. After editing, enter Ctrl-X to save the file and exit the editor. The available editor commands are displayed on screen.

Detailed information on using the Nano editor is available at this URL: <http://www.nano-editor.org/>.

Not all Nano editor features are supported on Cisco IOS XR software.

### Editing Routing Policy Configuration Elements Using the Emacs Editor

To edit the contents of a routing policy using the Emacs editor, use the following CLI command in EXEC mode:

```
edit route-policy name emacs
```

A copy of the route policy is copied to a temporary file and the editor is launched. After editing, save the editor buffer by using the Ctrl-X and Ctrl-S keystrokes. To save and exit the editor, use the Ctrl-X and Ctrl-C keystrokes. When you quit the editor, the buffer is committed. If there are no parse errors, the configuration is committed:

```
RP/0/RP0/CPU0:router# edit route-policy policy_A
-----
== MicroEMACS 3.8b () == rpl_edit.139281 ==
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

== MicroEMACS 3.8b () == rpl_edit.139281 ==
Parsing.
83 bytes parsed in 1 sec (82)bytes/sec
Committing.
1 items committed in 1 sec (0)items/sec
Updating.
Updated Commit database in 1 sec

RP/0/RP0/CPU0:router#
```

If there are parse errors, you are asked whether editing should continue:

```
RP/0/RP0/CPU0:router#edit route-policy policy_B
== MicroEMACS 3.8b () == rpl_edit.141738
route-policy policy_B
  set metric-type type_1
  if destination in (2001::/8) then
    drop
  endif
end-policy
!
== MicroEMACS 3.8b () == rpl_edit.141738 ==
Parsing.
105 bytes parsed in 1 sec (103)bytes/sec

% Syntax/Authorization errors in one or more commands.!! CONFIGURATION
FAILED DUE TO SYNTAX/AUTHORIZATION ERRORS
  set metric-type type_1
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

Continue editing? [no]:
```

If you answer **yes**, the editor continues on the text buffer from where you left off. If you answer **no**, the running configuration is not changed and the editing session is ended.

### Editing Routing Policy Configuration Elements Using the Vim Editor

Editing elements of a routing policy with Vim (Vi IMproved) is similar to editing them with Emacs except for some feature differences such as the keystrokes to save and quit. To write to a current file and exit, use the **:wq** or **:x** or **ZZ** keystrokes. To quit and confirm, use the **:q** keystrokes. To quit and discard changes, use the **:q!** keystrokes.

You can reference detailed online documentation for Vim at this URL: <http://www.vim.org/>

### Editing Routing Policy Configuration Elements Using the CLI

The CLI allows you to enter and delete route policy statements. You can complete a policy configuration block by entering applicable commands such as **end-policy** or **end-set**. Alternatively, the CLI interpreter allows you to use the **exit** command to complete a policy configuration block. The **abort** command is used to discard the current policy configuration and return to global configuration mode.

## How to Implement Routing Policy on Cisco IOS XR Software

This section contains the following procedures:

- [Defining a Route Policy, page RC-409](#) (required)
- [Attaching a Routing Policy to a BGP Neighbor, page RC-410](#) (required)
- [Modifying a Routing Policy Using a Text Editor, page RC-412](#) (optional)

## Defining a Route Policy

This task explains how to define a route policy.



### Note

If you want to modify an existing routing policy using the command-line interface (CLI), you must redefine the policy by completing this task.

### SUMMARY STEPS

1. **configure**
2. **route-policy** *name* [*parameter1*, *parameter2*, . . . , *parameterN*]
3. **end-policy**
4. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>route-policy</b> <i>name</i> [ <i>parameter1</i> , <i>parameter2</i> , . . . , <i>parameterN</i> ]  <b>Example:</b> RP/0/RP0/CPU0:router(config)# route-policy sample1	Enters route-policy configuration mode. <ul style="list-style-type: none"> <li>• After the route-policy has been entered, a group of commands can be entered to define the route-policy.</li> </ul>

	Command or Action	Purpose
Step 3	<code>end-policy</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config-rpl)# end-policy	Ends the definition of a route policy and exits route-policy configuration mode.
Step 4	<code>end</code> OR <code>commit</code>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# end OR RP/0/RP0/CPU0:router(config)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul> <p>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</p>

## Attaching a Routing Policy to a BGP Neighbor

This task explains how to attach a routing policy to a BGP neighbor.

### Prerequisites

A routing policy must be preconfigured and well defined prior to it being applied at an attach point. If a policy is not predefined, an error message is generated stating that the policy is not defined.

### SUMMARY STEPS

- `configure`
- `router bgp as-number`
- `neighbor ip-address`
- `address-family {ipv4 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | ipv4 mdt | ipv6 unicast | ipv6 multicast | ipv6 labeled-unicast | vpnv4 unicast | vpnv6 unicast}`
- `route-policy route-policy-name {in | out}`
- `end`  
OR  
`commit`



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router bgp 125</p>	<p>Configures a BGP routing process and enters router configuration mode.</p> <ul style="list-style-type: none"> <li>The <i>as-number</i> argument identifies the autonomous system in which the router resides. Valid values are from 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535.</li> </ul>
Step 3	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.20</p>	Specifies a neighbor IP address.
Step 4	<p><b>address-family</b> {<i>ipv4 unicast</i>   <i>ipv4 multicast</i>   <i>ipv4 labeled-unicast</i>   <i>ipv4 tunnel</i>   <i>ipv4 mdt</i>   <i>ipv6 unicast</i>   <i>ipv6 multicast</i>   <i>ipv6 labeled-unicast</i>   <i>vpn4 unicast</i>   <i>vpn6 unicast</i>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast</p>	Specifies the address family.

	Command or Action	Purpose
Step 5	<pre>route-policy policy-name {in   out}</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy example1 in </p>	Attaches the route-policy, which must be well formed and predefined.
Step 6	<pre>end</pre> <p>or</p> <pre>commit</pre> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-bgp-nbr-af)# end or RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</li> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <p>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</p>

## Modifying a Routing Policy Using a Text Editor

This task explains how to modify an existing routing policy using a text editor. See the “[Editing Routing Policy Configuration Elements](#)” section on page [RC-407](#) for information on text editors.

### SUMMARY STEPS

- edit** {route-policy | prefix-set | as-path-set | community-set | extcommunity-set {rt | soo} | policy-global | rd-set} name [nano | emacs | vim | inline {add | prepend | remove} set-element]
- show rpl route-policy** [name [detail] | states | brief]
- show rpl prefix-set** [name | states | brief]

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>edit {route-policy   prefix-set   as-path-set   community-set   extcommunity-set {rt   soo}   policy-global   rd-set} name [nano   emacs   vim   inline {add   prepend   remove} set-element]</pre> <p>Example: RP/0/RP0/CPU0:router# edit route-policy sample1</p>	<p>Identifies the route policy, prefix set, AS path set, community set, or extended community set name to be modified.</p> <ul style="list-style-type: none"> <li>A copy of the route policy, prefix set, AS path set, community set, or extended community set is copied to a temporary file and the editor is launched.</li> <li>After editing with Nano, save the editor buffer and exit the editor by using the Ctrl-X keystroke.</li> <li>After editing with Emacs, save the editor buffer by using the Ctrl-X and Ctrl-S keystrokes. To save and exit the editor, use the Ctrl-X and Ctrl-C keystrokes.</li> <li>After editing with Vim, to write to a current file and exit use the :wq or :x or ZZ keystrokes. To quit and confirm, use the :q keystrokes. To quit and discard changes, use the :q! keystrokes.</li> </ul>
Step 2	<pre>show rpl route-policy [name [detail]   states   brief]</pre> <p>Example: RP/0/RP0/CPU0:router# show rpl route-policy sample2</p>	<p>(Optional) Displays the configuration of a specific named route policy.</p> <ul style="list-style-type: none"> <li>Use the <b>detail</b> keyword to display all policies and sets that a policy uses.</li> <li>Use the <b>states</b> keyword to display all unused, inactive, and active states.</li> <li>Use the <b>brief</b> keyword to list the names of all extended community sets without their configurations.</li> </ul>
Step 3	<pre>show rpl prefix-set [name   states   brief]</pre> <p>Example: RP/0/RP0/CPU0:router# show rpl prefix-set prefixset1</p>	<p>(Optional) Displays the contents of a named prefix set.</p> <ul style="list-style-type: none"> <li>To display the contents of a named AS path set, community set, or extended community set, replace the <b>prefix-set</b> keyword with <b>as-path-set</b>, <b>community-set</b>, or <b>extcommunity-set</b>, respectively.</li> </ul>

## Configuration Examples for Implementing Routing Policy on Cisco IOS XR Software

This section provides the following configuration examples:

- [Routing Policy Definition: Example, page RC-414](#)
- [Simple Inbound Policy: Example, page RC-414](#)
- [Modular Inbound Policy: Example, page RC-415](#)
- [Translating Cisco IOS Route Maps to Cisco IOS XR Routing Policy Language: Example, page RC-416](#)

## Routing Policy Definition: Example

In the following example, a BGP route policy named `sample1` is defined using the `route-policy name` command. The policy compares the network layer reachability information (NLRI) to the elements in the prefix set `test`. If it evaluates to true, the policy performs the operations in the `then` clause. If it evaluates to false, the policy performs the operations in the `else` clause, that is, sets the MED value to 200 and adds the community 2:100 to the route. The final steps of the example commit the configuration to the router, exit configuration mode, and display the contents of route policy `sample1`.

```
configure
  route-policy sample1
    if destination in test then
      drop
    else
      set med 200
      set community (2:100) additive
    endif
  end-policy
end
show config running route-policy sample1
```

```
Building configuration...
route-policy sample1
  if destination in test then
    drop
  else
    set med 200
    set community (2:100) additive
  endif
end-policy
```

## Simple Inbound Policy: Example

The following policy discards any route whose network layer reachability information (NLRI) specifies a prefix longer than /24, and any route whose NLRI specifies a destination in the address space reserved by RFC 1918. For all remaining routes, it sets the MED and local preference, and adds a community to the list in the route.

For routes whose community lists include any values in the range from 101:202 to 106:202 that have a 16-bit tag portion containing the value 202, the policy prepends autonomous system number 2 twice, and adds the community 2:666 to the list in the route. Of these routes, if the MED is either 666 or 225, then the policy sets the origin of the route to incomplete, and otherwise sets the origin to IGP.

For routes whose community lists do not include any of the values in the range from 101:202 to 106:202, the policy adds the community 2:999 to the list in the route.

```
prefix-set too-specific
  0.0.0.0/0 ge 25 le 32
end-set

prefix-set rfc1918
  10.0.0.0/8 le 32,
  172.16.0.0/12 le 32,
  192.168.0.0/16 le 32
end-set

route-policy inbound-tx
  if destination in too-specific or destination in rfc1918 then
    drop
```

```

endif
set med 1000
set local-preference 90
set community (2:1001) additive
if community matches-any ([101..106]:202) then
  prepend as-path 2.30 2
  set community (2:666) additive
  if med is 666 or med is 225 then
    set origin incomplete
  else
    set origin igp
  endif
else
  set community (2:999) additive
endif
end-policy

router bgp 2
neighbor 10.0.1.2 address-family ipv4 unicast route-policy inbound-tx in

```

## Modular Inbound Policy: Example

The following policy example shows how to build two inbound policies, in-100 and in-101, for two different peers. In building the specific policies for those peers, the policy reuses some common blocks of policy that may be common to multiple peers. It builds a few basic building blocks, the policies common-inbound, filter-bogons, and set-lpref-prepend.

The filter-bogons building block is a simple policy that filters all undesirable routes, such as those from the RFC 1918 address space. The policy set-lpref-prepend is a utility policy that can set the local preference and prepend the AS path according to parameterized values that are passed in. The common-inbound policy uses these filter-bogons building blocks to build a common block of inbound policy. The common-inbound policy is used as a building block in the construction of in-100 and in-101 along with the set-lpref-prepend building block.

This is a simple example that illustrates the modular capabilities of the policy language.

```

prefix-set bogon
  10.0.0.0/8 ge 8 le 32,
  0.0.0.0,
  0.0.0.0/0 ge 27 le 32,
  192.168.0.0/16 ge 16 le 32
end-set
!
route-policy in-100
  apply common-inbound
  if community matches-any ([100..120]:135) then
    apply set-lpref-prepend (100,100,2)
    set community (2:1234) additive
  else
    set local-preference 110
  endif
  if community matches-any ([100..666]:[100..999]) then
    set med 444
    set local-preference 200
    set community (no-export) additive
  endif
end-policy
!
route-policy in-101
  apply common-inbound
  if community matches-any ([101..200]:201) then

```

```

        apply set-lpref-prepend(100,101,2)
        set community (2:1234) additive
    else
        set local-preference 125
    endif
end-policy
!
route-policy filter-bogons
    if destination in bogon then
        drop
    else
        pass
    endif
end-policy
!
route-policy common-inbound
    apply filter-bogons
    set origin igp
    set community (2:333)
end-policy
!
route-policy set-lpref-prepend($lpref,$as,$prependcnt)
    set local-preference $lpref
    prepend as-path $as $prependcnt
end-policy

```

## Translating Cisco IOS Route Maps to Cisco IOS XR Routing Policy Language: Example

RPL performs the same functions as route-maps. See *Converting Cisco IOS Configurations to Cisco IOS XR Configurations*.

## Additional References

The following sections provide references related to implementing RPL.

## Related Documents

Related Topic	Document Title
Routing policy language commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Routing Policy Language Commands on Cisco IOS XR</i>
Regular expression syntax	“Understanding Regular Expressions, Special Characters and Patterns” appendix in the <i>Cisco IOS XR Getting Started Guide</i>

## Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
RFC 1771	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 4360	BGP Extended Communities Attribute

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>







# Implementing Static Routes on Cisco IOS XR Software

*Static routes* are user-defined routes that cause packets moving between a source and a destination to take a specified path. Static routes can be important if the Cisco IOS XR software cannot build a route to a particular destination. They are useful for specifying a gateway of last resort to which all unroutable packets are sent.

This module describes the tasks you need to implement static routes on your Cisco IOS XR network.



## Note

For more information about static routes on the Cisco IOS XR software and complete descriptions of the static routes commands listed in this module, see the “[Related Documents](#)” section of this module. To locate documentation for other commands that might appear while executing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing Static Routes on Cisco IOS XR Software

Release	Modification
Release 2.0	This feature was introduced on the Cisco CRS-1.
Release 3.0	No modification.
Release 3.2	Support was added for the Cisco XR 12000 Series Router.
Release 3.3.0	Support for configuring static routes in static router configuration mode was added.  The <b>route</b> command was replaced with the <b>router static</b> command.  The <b>route maximum</b> command was replaced with the <b>maximum path</b> command.  VPN routing and forwarding (VRF) support was added to the command syntax.
Release 3.4.0	No modification.
Release 3.5.0	IPv6 Provider Edge and IPv6 VPN Provider Edge over Multiprotocol Label Switching support was added.
Release 3.6.0	No modification.

# Contents

- [Prerequisites for Implementing Static Routes on Cisco IOS XR Software, page RC-420](#)
- [Information About Implementing Static Routes on Cisco IOS XR Software, page RC-420](#)
- [How to Implement Static Routes on Cisco IOS XR Software, page RC-424](#)
- [Configuration Examples, page RC-432](#)
- [Where to Go Next, page RC-433](#)
- [Additional References, page RC-433](#)

## Prerequisites for Implementing Static Routes on Cisco IOS XR Software

You must be in a user group associated with a task group that includes the proper task IDs for static routing commands. For detailed information about user groups and task IDs, see the *Configuring AAA Services on Cisco IOS XR Software* module of *Cisco IOS XR System Security Configuration Guide*.

## Information About Implementing Static Routes on Cisco IOS XR Software

To implement static routes you need to understand the following concepts:

- [Static Route Functional Overview, page RC-420](#)
- [Default Administrative Distance, page RC-421](#)
- [Directly Connected Routes, page RC-421](#)
- [Recursive Static Routes, page RC-422](#)
- [Fully Specified Static Routes, page RC-422](#)
- [Floating Static Routes, page RC-423](#)
- [Default VRF, page RC-423](#)
- [IPv4 and IPv6 Static VRF Routes, page RC-423](#)
- [IPv6/IPv6 VPN Provider Edge Transport over MPLS, page RC-423](#)

## Static Route Functional Overview

Static routes are entirely user configurable and can point to a next-hop interface, next-hop IP address, or both. In Cisco IOS XR software, if an interface was specified, then the static route is installed in the Routing Information Base (RIB) if the interface is reachable. If an interface was not specified, the route is installed if the next-hop address is reachable. The only exception to this configuration is when a static route is configured with the permanent attribute, in which case it is installed in RIB regardless of reachability.

Networking devices forward packets using route information that is either manually configured or dynamically learned using a routing protocol. Static routes are manually configured and define an explicit path between two networking devices. Unlike a dynamic routing protocol, static routes are not automatically updated and must be manually reconfigured if the network topology changes. The benefits of using static routes include security and resource efficiency. Static routes use less bandwidth than dynamic routing protocols, and no CPU cycles are used to calculate and communicate routes. The main disadvantage to using static routes is the lack of automatic reconfiguration if the network topology changes.

Static routes can be redistributed into dynamic routing protocols, but routes generated by dynamic routing protocols cannot be redistributed into the static routing table. No algorithm exists to prevent the configuration of routing loops that use static routes.

Static routes are useful for smaller networks with only one path to an outside network and to provide security for a larger network for certain types of traffic or links to other networks that need more control. In general, most networks use dynamic routing protocols to communicate between networking devices but may have one or two static routes configured for special cases.

**Note**

---

For information on configuring static routes to distribute Multiprotocol Label Switching (MPLS) Layer 3 virtual private network (VPN) information, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

---

## Default Administrative Distance

Static routes have a default administrative distance of 1. A low number indicates a preferred route. By default, static routes are preferred to routes learned by routing protocols. Therefore, you can configure an administrative distance with a static route if you want the static route to be overridden by dynamic routes. For example, you could have routes installed by the Open Shortest Path First (OSPF) protocol with an administrative distance of 120. To have a static route that would be overridden by an OSPF dynamic route, specify an administrative distance greater than 120.

## Directly Connected Routes

The routing table considers the static routes that point to an interface as “directly connected.” Directly connected networks are advertised by IGP routing protocols if a corresponding **interface** command is contained under the router configuration stanza of that protocol.

In directly attached static routes, only the output interface is specified. The destination is assumed to be directly attached to this interface, so the packet destination is used as the next hop address. The following example shows how to specify that all destinations with address prefix 2001:0DB8::/32 are directly reachable through interface GigabitEthernet 0/5/0/0:

```
RP/0/RP0/CPU0:router(config)# router static  
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast  
RP/0/RP0/CPU0:router(config-static-afi)# 2001:0DB8::/32 gigabitethernet 0/5/0/0
```

Directly attached static routes are candidates for insertion in the routing table only if they refer to a valid interface; that is, an interface that is both up and has IPv4 or IPv6 enabled on it.

## Recursive Static Routes

In a recursive static route, only the next hop is specified. The output interface is derived from the next hop. The following example shows how to specify that all destinations with address prefix 2001:0DB8::/32 are reachable through the host with address 2001:0DB8:3000::1:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 2001:0DB8::/32 2001:0DB8:3000::1
```

A recursive static route is valid (that is, it is a candidate for insertion in the routing table) only when the specified next hop resolves, either directly or indirectly, to a valid output interface, provided the route does not self-recurse, and the recursion depth does not exceed the maximum IPv6 forwarding recursion depth.

A route self-recurses if it is itself used to resolve its own next hop. If a static route becomes self-recursive, RIB sends a notification to static routes to withdraw the recursive route.

Assuming a BGP route 2001:0DB8:3000::0/16 with next hop of 2001:0DB8::0104, the following static route would not be inserted into the IPv6 RIB because the BGP route next hop resolves through the static route and the static route resolves through the BGP route making it self-recursive:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 001:0DB8::/32 2001:0DB8:3000::1
```

This static route is not inserted into the IPv6 routing table because it is self-recursive. The next hop of the static route, 2001:0DB8:3000:1, resolves through the BGP route 2001:0DB8:3000:0/16, which is itself a recursive route (that is, it only specifies a next hop). The next hop of the BGP route, 2001:0DB8::0104, resolves through the static route. Therefore, the static route would be used to resolve its own next hop.

It is not normally useful to manually configure a self-recursive static route, although it is not prohibited. However, a recursive static route that has been inserted in the routing table may become self-recursive as a result of some transient change in the network learned through a dynamic routing protocol. If this occurs, the fact that the static route has become self-recursive will be detected and it will be removed from the routing table, although not from the configuration. A subsequent network change may cause the static route to no longer be self-recursive, in which case it will be re-inserted in the routing table.

## Fully Specified Static Routes

In a fully specified static route, both the output interface and next hop are specified. This form of static route is used when the output interface is multiaccess and it is necessary to explicitly identify the next hop. The next hop must be directly attached to the specified output interface. The following example shows a definition of a fully specified static route:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 2001:0DB8::/32 Gigetherne0/0/0/0
2001:0DB8:3000::1
```

A fully specified route is valid (that is, a candidate for insertion into the routing table) when the specified interface, IPv4 or IPv6, is enabled and up.

## Floating Static Routes

Floating static routes are static routes that are used to back up dynamic routes learned through configured routing protocols. A floating static route is configured with a higher administrative distance than the dynamic routing protocol it is backing up. As a result, the dynamic route learned through the routing protocol is always preferred to the floating static route. If the dynamic route learned through the routing protocol is lost, the floating static route is used in its place. The following example shows how to define a floating static route:

```
RP/0/RP0/CPU0:router(config)# router static  
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast  
RP/0/RP0/CPU0:router(config-static-afi)# 2001:0DB8::/32 2001:0DB8:3000::1 210
```

Any of the three types of static routes can be used as a floating static route. A floating static route must be configured with an administrative distance that is greater than the administrative distance of the dynamic routing protocol because routes with smaller administrative distances are preferred.

**Note**

---

By default, static routes have smaller administrative distances than dynamic routes, so static routes are preferred to dynamic routes.

---

## Default VRF

A static route is always associated with a VPN routing and forwarding (VRF) instance. The VRF can be the default VRF or a specified VRF. Specifying a VRF, using the **vrf vrf-name** command, allows you to enter VRF configuration mode for a specific VRF where you can configure a static route. If a VRF is not specified, a default VRF static route is configured.

## IPv4 and IPv6 Static VRF Routes

An IPv4 or IPv6 static VRF route is the same as a static route configured for the default VRF. The IPv4 and IPv6 address families are supported in each VRF.

## IPv6/IPv6 VPN Provider Edge Transport over MPLS

IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) leverages the existing Multiprotocol Label Switching (MPLS) IPv4 core infrastructure for IPv6 transport. 6PE and 6VPE enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs).

**Note**

---

This feature is supported on Cisco XR 12000 Series Routers.

---

Static routes can be configured under the default VRF for 6PE functionality and under IPv6 VPN routing and forwarding (VRF) instances for 6VPE functionality.

For detailed information about configuring 6PE and 6VPE over MPLS, see *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*.

# How to Implement Static Routes on Cisco IOS XR Software

This section contains the following procedures:

- [Configuring a Static Route, page RC-424](#) (required)
- [Configuring a Floating Static Route, page RC-425](#) (optional)
- [Configuring Static Routes Between PE-CE Routers, page RC-427](#) (optional)
- [Changing the Maximum Number of Allowable Static Routes, page RC-429](#) (optional)
- [Associating a VRF with a Static Route, page RC-430](#) (optional)

## Configuring a Static Route

This task explains how to configure a static route.

### SUMMARY STEPS

1. **configure**
2. **router static**
3. **vrf vrf-name**
4. **address-family {ipv4 | ipv6} {unicast | multicast}**
5. *prefix mask [vrf vrf-name] {ip-address | interface-type interface-instance} [distance] [description text] [tag tag] [permanent]*
6. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router static</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router static	Enters static route configuration mode.
Step 3	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static)# vrf vrf_A	(Optional) Enters VRF configuration mode. If a VRF is not specified, the static route is configured under the default VRF.

	Command or Action	Purpose
Step 4	<p><b>address-family</b> {<b>ipv4</b>   <b>ipv6</b>} {<b>unicast</b>   <b>multicast</b>}</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf)# address family ipv4 unicast</p>	Enters address family mode.
Step 5	<p><b>prefix mask</b> [<b>vrf vrf-name</b>] {<b>ip-address</b>   <b>interface-type interface-instance</b>} [<b>distance</b>] [<b>description text</b>] [<b>tag tag</b>] [<b>permanent</b>]</p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# 10.0.0.0/8 172.20.16.6 110</p>	<p>Configures an administrative distance of 110.</p> <ul style="list-style-type: none"> <li>This example shows how to route packets for network 10.0.0.0 through to a next hop at 172.20.16.6 if dynamic information with administrative distance less than 110 is not available.</li> </ul>
Step 6	<p><b>end</b> or <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# end or RP/0/RP0/CPU0:router(config-static-vrf-afi)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a Floating Static Route

This task explains how to configure a floating static route.

### SUMMARY STEPS

1. **configure**
2. **router static**
3. **vrf vrf-name**
4. **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}
5. **prefix mask** [**vrf vrf-name**] {**ip-address** | **interface-type interface-instance**} [**distance**] [**description text**] [**tag tag**] [**permanent**]

```

6. end
   or
   commit

```

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router static</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router static	Enters static route configuration mode.
Step 3	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static)# vrf vrf_A	(Optional) Enters VRF configuration mode.  If a VRF is not specified, the static route is configured under the default VRF.
Step 4	<b>address-family {ipv4   ipv6} {unicast   multicast}</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf)# address family ipv6 unicast	Enters address family mode.



	Command or Action	Purpose
Step 5	<pre>prefix mask [vrf vrf-name] {ip-address   interface-type interface-instance} [distance] [description text] [tag tag] [permanent]</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-static-vrf-afi)# 20 01:0DB8::/32 2001:0DB8:3000::1 201 </p>	Configures an administrative distance of 201.
Step 6	<pre>end OR commit</pre> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-static-vrf-afi)# en d OR RP/0/RP0/CPU0:router(config-static-vrf-afi)# commit </p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  Uncommitted changes found, commit them before exiting (yes/no/cancel)?  [cancel]:</li> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Static Routes Between PE-CE Routers

This task explains how to configure static routing between PE-CE routers.



### Note

VRF fallback is not supported with IPv6 VPN Provider Edge (6VPE).

### SUMMARY STEPS

- configure
- router static
- vrf vrf-name
- address-family {ipv4 | ipv6} {unicast | multicast}
- prefix mask [vrf vrf-name] {ip-address | interface-type interface-instance} [distance] [description text] [tag tag] [permanent]
- end  
or  
commit

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router static</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router static	Enters static route configuration mode.
Step 3	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static)# vrf vrf_A	(Optional) Enters VRF configuration mode. If a VRF is not specified, the static route is configured under the default VRF.
Step 4	<b>address-family {ipv4   ipv6} {unicast   multicast}</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf)# address family ipv6 unicast	Enters address family mode.
Step 5	<b>prefix mask [vrf vrf-name] {ip-address   interface-type interface-instance} [distance] [description text] [tag tag] [permanent]</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# 2001:0DB8::/32 2001:0DB8:3000::1 201	Configures an administrative distance of 201.
Step 6	<b>end</b> OR <b>commit</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# end OR RP/0/RP0/CPU0:router(config-static-vrf-afi)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Changing the Maximum Number of Allowable Static Routes

This task explains how to change the maximum number of allowable static routes.

### Restrictions

The number of static routes that can be configured on a router for a given address family is limited by default to 4000. The limit can be raised or lowered using the **maximum path** command. Note that if you use the **maximum path** command to reduce the configured maximum allowed number of static routes for a given address family below the number of static routes currently configured, the change is rejected. In addition, understand the following behavior: If you commit a batch of routes that would, when grouped, push the number of static routes configured above the maximum allowed, the first  $n$  routes in the batch are accepted. The number previously configured is accepted, and the remainder are rejected. The  $n$  argument is the difference between the maximum number allowed and number previously configured.

### SUMMARY STEPS

1. **configure**
2. **router static**
3. **maximum path {ipv4 | ipv6} value**
4. **end**  
or  
**commit**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router static</b>  <b>Example:</b> RP/0/RP0/CPU0:router(config)# router static	Enters static route configuration mode.

	Command or Action	Purpose
Step 3	<p><b>maximum path</b> {<b>ipv4</b>   <b>ipv6</b>} <i>value</i></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-static)# maximum path ipv4 10000</p>	<p>Changes the maximum number of allowable static routes.</p> <ul style="list-style-type: none"> <li>Specify IPv4 or IPv6 address prefixes.</li> <li>Specify the maximum number of static routes for the given address family. The range is from 1 to 140000.</li> <li>This example sets the maximum number of static IPv4 routes to 10000.</li> </ul>
Step 4	<p><b>end</b>  OR  <b>commit</b></p> <p><b>Example:</b>  RP/0/RP0/CPU0:router(config-static)# end  OR  RP/0/RP0/CPU0:router(config-static)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes:  <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Associating a VRF with a Static Route

This task explains how to associate a VRF with a static route.

### SUMMARY STEPS

- configure**
- router static**
- vrf** *vrf-name*
- address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}
- prefix mask* [**vrf** *vrf-name*] {*ip-address* | *interface-type interface-instance*} [*distance*] [**description** *text*] [**tag** *tag*] [**permanent**]
- end**  
or  
**commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router# configure</p>	Enters global configuration mode.
Step 2	<p><b>router static</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config)# router static</p>	Enters static route configuration mode.
Step 3	<p><b>vrf vrf-name</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static)# vrf vrf_A</p>	Enters VRF configuration mode.
Step 4	<p><b>address-family {ipv4   ipv6} {unicast   multicast}</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf)# address family ipv6 unicast</p>	Enters address family mode.
Step 5	<p><b>prefix mask [vrf vrf-name] {ip-address   interface-type interface-instance} [distance] [description text] [tag tag] [permanent]</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# 20 01:0DB8::/32 2001:0DB8:3000::1 201</p>	Configures an administrative distance of 201.
Step 6	<p><b>end</b> OR <b>commit</b></p> <p><b>Example:</b> RP/0/RP0/CPU0:router(config-static-vrf-afi)# end OR RP/0/RP0/CPU0:router(config-static-vrf-afi)# commit</p>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>When you issue the <b>end</b> command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> <li>Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

# Configuration Examples

This section provides the following configuration examples:

- [Configuring Traffic Discard: Example](#)
- [Configuring a Fixed Default Route: Example](#)
- [Configuring a Floating Static Route: Example](#)
- [Configuring a Static Route Between PE-CE Routers: Example](#)

## Configuring Traffic Discard: Example

Configuring a static route to point at interface null 0 may be used for discarding traffic to a particular prefix. For example, if it is required to discard all traffic to prefix 2001:0DB8:42:1/64, the following static route would be defined:

```
configure
router static
  address-family ipv6 unicast
    2001:0DB8:42:1::/64 null 0
  end
```

## Configuring a Fixed Default Route: Example

A default static route is often used in simple router topologies. In the following example, a route is configured with an administrative distance of 110.

```
configure
router static
  address-family ipv4 unicast
    0.0.0.0/0 2.6.0.1 110
  end
```

## Configuring a Floating Static Route: Example

A floating static route is often used to provide a backup path if connectivity fails. In the following example, a route is configured with an administrative distance of 201.

```
configure
router static
  address-family ipv6 unicast
    2001:0DB8::/32 2001:0DB8:3000::1 201
  end
```

## Configuring a Static Route Between PE-CE Routers: Example

In the following example, a static route between PE and CE routers is configured, and a VRF is associated with the static route:

```
configure
router static
  vrf vrf_A
  address-family ipv4 unicast
```

```
0.0.0.0/0 2.6.0.2 120
end
```

## Where to Go Next

For additional information about static routes, routing protocols, and RIB, consult the following publications:

- *Implementing and Monitoring RIB on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *Implementing BGP on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *Implementing EIGRP on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *Implementing IS-IS on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *Implementing MPLS Layer 3 VPN on Cisco IOS XR Software* in *Cisco IOS XR Multiprotocol Label Switching Configuration Guide*
- *Implementing OSPF on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *Implementing OSPFv3 on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*
- *RIB Commands on Cisco IOS XR Software* in *Cisco IOS XR Routing Command Reference*
- *Implementing RIP on Cisco IOS XR Software* in *Cisco IOS XR Routing Configuration Guide*

## Additional References

The following sections provide references related to implementing static routes on Cisco IOS XR software.

## Related Documents

Related Topic	Document Title
Static routes commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Static Routing Commands on Cisco IOS XR Software</i> , Release 3.6
MPLS Layer 3 VPN configuration: configuration concepts, task, and examples	<i>Cisco IOS XR Multiprotocol Label Switching Configuration Guide</i> , Release 3.6

## Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

## MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

## Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>





## INDEX

---

### A

address family command [RC-8](#), [RC-268](#)  
address-family command (IS-IS) [RC-283](#)  
address-family ipv4 command [RC-149](#)  
address-family ipv4 mdt command [RC-116](#)  
adjacencies, tuning [RC-298](#)  
administrative distance [RC-321](#)  
administrative distance, static routes [RC-421](#)  
aggregate-address command [RC-61](#), [RC-107](#)  
allowas-in command [RC-108](#)  
apply command [RC-374](#)  
Area Border Routers (ABRs) [RC-178](#)  
area command [RC-196](#)  
as-override command [RC-108](#)  
attached bit on an IS-IS instance [RC-273](#)  
authentication  
    MD5 (OSPFv2) [RC-180](#)  
    route, key rollover (OSPFv2) [RC-180](#)  
    strategies [RC-180](#)  
authentication, configuring (OSPFv2) [RC-204](#)  
authentication command (OSPFv2) [RC-205](#)  
authentication message-digest command [RC-211](#)  
Autonomous System Boundary Routers (ASBRs) [RC-178](#)  
autonomous system number format [RC-7](#)  
autonomous systems [RC-177](#)  
auto-summary command [RC-340](#)

---

### B

backbone area [RC-177](#)  
bandwidth-percent command [RC-151](#)  
bestpath algorithm [RC-27](#)

### BGP (Border Gateway Protocol)

autonomous system number format [RC-7](#)  
bestpath algorithm [RC-27](#)  
BGP keychains [RC-39](#)  
bidirectional forwarding detection [RC-3](#)  
configuration  
    grouping [RC-9](#)  
    inheriting [RC-11](#)  
default address family [RC-37](#)  
description [RC-1](#)  
distributed BGP [RC-37](#)  
functional overview [RC-3](#)  
inheritance, monitoring [RC-15](#)  
IPv6 and IPv6 VPN provider edge transport over MPLS [RC-39](#)  
IPv6 provider edge multipath [RC-40](#)  
MPLS VPN carrier supporting carrier [RC-38](#)  
multicast VPN [RC-41](#)  
multiprotocol [RC-32](#)  
neighbors, maximum limits on [RC-4](#)  
policy attach points  
    allocate label [RC-383](#)  
    clear policy [RC-384](#)  
    dampening [RC-376](#)  
    debug [RC-384](#)  
    default originate [RC-377](#)  
    export [RC-381](#)  
    import [RC-381](#)  
    neighbor export [RC-377](#)  
    neighbor import [RC-378](#)  
    neighbor-orf [RC-383](#)  
    network [RC-378](#)  
    next-hop [RC-384](#)

- redistribute [RC-379](#)
- retain route target [RC-382](#)
- show bgp [RC-379](#)
- table policy [RC-380](#)
- policy attach points, aggregation [RC-375](#)
- router identifier [RC-3](#)
- routing policy
  - enforcing [RC-20](#)
- update groups
  - description [RC-22](#)
  - example [RC-129](#)
- VPNv4 and VPNv6 over the IP core using L2TPv3 tunnels [RC-40](#)
- bgp address family submode
  - bgp client-to-client reflection disable command [RC-102](#)
  - bgp dampening command [RC-101](#)
  - retain route-target command [RC-105](#)
- bgp bestpath as-path ignore command [RC-57](#)
- bgp bestpath compare-routerid command [RC-58](#)
- bgp bestpath med always command [RC-57](#)
- bgp bestpath med confed command [RC-57](#)
- bgp bestpath med missing-as-worst command [RC-57](#)
- bgp client-to-client reflection disable command [RC-102](#)
- bgp confederation identifier command [RC-48](#)
- bgp confederation peers command [RC-49](#)
- bgp cost community [RC-22](#)
- bgp dampening command [RC-67](#), [RC-101](#)
- bgp default local-preference command [RC-52](#)
- bgp global address family submode [RC-8](#)
  - aggregate-address command [RC-61](#)
  - bgp dampening command [RC-67](#)
  - distance bgp command [RC-74](#)
  - network command [RC-59](#)
  - redistribute command [RC-65](#)
  - See* address family command
  - table-policy command [RC-72](#)
- bgp neighbor address family submode [RC-8](#), [RC-269](#)
  - next-hop self command [RC-84](#)
  - route-policy (BGP) command [RC-81](#)
  - route-policy command [RC-47](#), [RC-76](#), [RC-103](#)
  - route-reflector-client command [RC-79](#)
  - See* neighbor address family command
  - send-community-ebgp command [RC-86](#)
  - soft-reconfiguration inbound always command [RC-93](#)
  - weight command [RC-56](#)
- bgp neighbor command [RC-9](#)
- bgp neighbor group submode [RC-10](#), [RC-76](#)
  - See* neighbor-group command
- bgp neighbor submode [RC-8](#), [RC-9](#), [RC-268](#)
  - keychain command [RC-115](#)
  - password command [RC-102](#)
  - route-policy command [RC-117](#)
  - See* bgp neighbor command
  - shutdown command [RC-102](#), [RC-119](#)
  - timers command [RC-102](#)
  - update-source command [RC-102](#)
  - use command [RC-76](#)
- bgp redistribute-internal command [RC-62](#)
- bgp router-id command [RC-99](#), [RC-106](#)
- bgp router submode [RC-8](#), [RC-268](#)
  - address-family ipv4 mdt command [RC-116](#)
  - bgp bestpath as-path ignore command [RC-57](#)
  - bgp bestpath compare-routerid command [RC-58](#)
  - bgp bestpath med always command [RC-57](#)
  - bgp bestpath med confed command [RC-57](#)
  - bgp bestpath med missing-as-worst command [RC-57](#)
  - bgp confederation identifier command [RC-48](#)
  - bgp confederation peers command [RC-49](#)
  - bgp default local-preference command [RC-52](#)
  - bgp redistribute-internal command [RC-62](#)
  - bgp router-id command [RC-99](#)
  - default-metric command [RC-54](#)
  - See* router bgp command
  - timers bgp command [RC-50](#)
- bgp session group submode [RC-10](#)
- BGP update groups example [RC-129](#)
- bgp VPNv4 address family submode [RC-9](#)

bgp VPNv6 address family submode [RC-9](#)  
 BGP vrf address family submode  
     export route-policy command [RC-97](#)  
     import route-policy command [RC-97](#)  
     import route-target command [RC-97](#)  
     maximum prefix command [RC-97](#)  
 bgp vrf address family submode [RC-8](#)  
     aggregate-address command [RC-107](#)  
     maximum-paths command [RC-110](#)  
     redistribute command [RC-113](#)  
 BGP VRF neighbor address family submode  
     as-override command [RC-108](#)  
     route-policy command [RC-108](#)  
     site-of-origin command [RC-108](#)  
 bgp VRF neighbor address family submode [RC-9](#)  
 bgp vrf neighbor address family submode  
     allowas-in command [RC-108](#)  
 BGP VRF neighbor submode  
     ebgp-multihop command [RC-107](#)  
 bgp VRF neighbor submode [RC-8](#)  
 bgp vrf neighbor submode  
     password command [RC-107](#)  
 BGP vrf submode  
     rd command [RC-100](#)  
 bgp vrf submode [RC-8](#)  
     bgp router-id command [RC-106](#)  
     label-allocation-mode per-ce command [RC-106](#)  
 bidirectional forwarding detection [RC-3](#)  
 broadcast-for-v2 command [RC-338](#)

## C

circuit-type command [RC-283, RC-313, RC-314](#)  
 clear bgp flap-statistics command [RC-69](#)  
 clear bgp flap-statistics reexp command [RC-70](#)  
 clear bgp flap-statistics route-policy command [RC-70](#)  
 clear bgp soft in command [RC-120](#)  
 clear bgp soft out command [RC-121](#)  
 clear eigrp neighbors command [RC-165](#)

clear eigrp topology command [RC-165](#)  
 clear ospf command [RC-255](#)  
 clear ospfv3 command [RC-255](#)  
 configuring  
     OSPFv3  
         graceful restart [RC-230](#)  
 cost community, BGP [RC-22](#)  
 csnp-interval command [RC-287](#)

## D

dampening, route [RC-33](#)  
 dead interval command [RC-201](#)  
 default address family [RC-19](#)  
 default-cost command [RC-198](#)  
 default-information originate command [RC-304](#)  
 default-metric command [RC-54, RC-150](#)  
 distance bgp command [RC-74](#)  
 distance command [RC-150, RC-305](#)  
 distributed BGP [RC-37](#)  
 draft-bonica-tcp-auth-05.txt [RC-133](#)  
 draft-ietf-idr-as4bytes-12.txt [RC-133](#)  
 draft-ietf-idr-avoid-transition-00.txt [RC-133](#)  
 draft-ietf-idr-bgp4-24.txt, BGP [RC-133](#)  
 draft-ietf-idr-bgp4-mib-15.txt, BGP [RC-133](#)  
 draft-ietf-idr-cease-subcode-05.txt [RC-133](#)  
 draft-ietf-isis-igp-p2p-over-lan-05.txt, Point-to-point  
 operation over LAN [RC-317](#)  
 draft-ietf-isis-ipv6-05.txt, Routing IPv6 with  
 IS-IS [RC-317](#)  
 draft-ietf-isis-restart-04.txt, Restart Signalling for  
 IS-IS [RC-317](#)  
 draft-ietf-isis-traffic-05.txt, IS-IS Extensions for Traffic  
 Engineering [RC-317](#)  
 draft-ietf-isis-wg-multi-topology-06.txt, M-ISIS  
     Multi Topology (MT) Routing in IS-IS [RC-317](#)  
 Draft-ietf-rtgwg-ipfrr-framework-06.txt, IP fast  
 reroute [RC-317, RC-331](#)  
 draft-ietf-rtgwg-lf-conv-frmwk-00.txt [RC-317, RC-331](#)  
 draft-nalawade-idr-mdt-safi-03.txt [RC-133](#)

**E**

ebgp-multihop command [RC-107](#)

edit command [RC-413](#)

EIGRP ()

policy attach points

default-accept-in [RC-401](#)

EIGRP (Enhanced Interior Gateway Routing Protocol)

features [RC-137](#)

hello interval and hold time [RC-145](#)

overview [RC-137](#)

policy attach points

default-accept-out [RC-401](#)

if-policy-in [RC-402](#)

if-policy-out [RC-402](#)

policy-in [RC-401](#)

policy-out [RC-402](#)

redistribute [RC-402](#)

restrictions [RC-158](#)

routing policy options [RC-146](#)

split horizon [RC-144](#)

stub routing [RC-145](#)

end-policy command [RC-45](#)

EXEC mode [RC-17, RC-166](#)

clear bgp flap statistics command [RC-69](#)

clear bgp flap statistics reexp command [RC-70](#)

clear bgp flap statistics route-policy command [RC-70](#)

clear bgp soft in command [RC-120](#)

clear bgp soft out command [RC-121](#)

clear eigrp neighbors command [RC-165](#)

clear eigrp topology command [RC-165](#)

clear ospf command [RC-255](#)

clear ospfv3 command [RC-255](#)

edit command [RC-413](#)

show bgp af-group command [RC-16, RC-17](#)

show bgp cidr-only command [RC-124, RC-125](#)

show bgp community command [RC-124, RC-125, RC-126](#)

show bgp count-only command [RC-125](#)

show bgp flap-statistics command [RC-68](#)

show bgp flap statistics reexp command [RC-68](#)

show bgp flap statistics route-policy command [RC-69](#)

show bgp inheritance command [RC-16](#)

show bgp neighbor command [RC-15](#)

show bgp neighbor-group command [RC-18, RC-125](#)

show bgp neighbors command [RC-124](#)

show bgp paths command [RC-124](#)

show bgp reexp command [RC-124](#)

show bgp summary command [RC-125](#)

show eigrp accounting command [RC-165](#)

show eigrp interfaces command [RC-165](#)

show eigrp neighbors command [RC-159](#)

show eigrp topology command [RC-166](#)

show eigrp traffic command [RC-166](#)

show isis command [RC-279](#)

show isis database command [RC-288](#)

show isis database-log command [RC-288](#)

show isis interface command [RC-301](#)

show isis lsp-log command [RC-288](#)

show isis mpls command [RC-297](#)

show isis mpls traffic-eng adjacency-log command [RC-297](#)

show isis mpls traffic-eng advertisements command [RC-297](#)

show isis neighbors command [RC-301](#)

show isis spf-log command [RC-303](#)

show isis topology command [RC-284](#)

show protocols eigrp command [RC-166](#)

show rpl route-policy command [RC-413](#)

show running-config command [RC-290](#)

export route-policy command [RC-97](#)

export route-target command [RC-98](#)

**F**

forwarding adjacency [RC-276](#)

---

**G**

Generalized TTL Security Mechanism (GTSM)  
     configuring virtual links [RC-193](#)

Generalized TTL Security Mechanism (GTSM), RFC 3682  
     TTL value [RC-193](#)

global parameters [RC-365](#)

graceful restart [RC-230](#)

graceful-restart helper command [RC-189](#)

graceful-restart interval command [RC-189](#)

graceful-restart lifetime command [RC-189](#)

---

**H**

hello-interval (IS-IS) command [RC-299](#)

hello interval (OSPF) command [RC-201](#)

hello-multiplier command [RC-299](#)

hello-padding command [RC-299](#)

hello-password command [RC-293, RC-295, RC-300](#)

holdtime command [RC-150](#)

---

**I**

if submode  
     ipv4 address command [RC-234, RC-239](#)

ignore-lsp-errors command [RC-287](#)

import route-policy command [RC-97](#)

import route-target command [RC-97](#)

inheritance  
     configurations (BGP) [RC-11](#)  
     monitoring [RC-15](#)

interarea tunnels [RC-276](#)

interface command [RC-150, RC-339](#)

interior routers [RC-179](#)

IP Fast Reroute [RC-323](#)

IP fast reroute [RC-276](#)  
     loop-free alternate [RC-317, RC-331](#)

ipv4 address command  
     10 Gigabit Ethernet [RC-234, RC-239](#)

## IPv6

IS-IS support  
     multitopology [RC-271](#)  
     single-topology [RC-271](#)

RIB support [RC-322](#)

routing [RC-269](#)

IPv6 and IPv6 VPN provider edge transport over  
 MPLS [RC-39, RC-322, RC-423](#)

IPv6 provider edge multipath [RC-40](#)

IS-IS (Intermediate System-to-Intermediate System)

adjacencies, tuning [RC-298](#)

attached bit on an instance [RC-273](#)

authentication, configuring [RC-291, RC-293](#)

Cisco IOS and Cisco IOS XR software differences,  
 configuration

    grouped [RC-268](#)

configuration

    grouped configuration [RC-268](#)

    Level 1 or Level 2 routing [RC-277](#)

    multitopology [RC-284](#)

    restrictions [RC-266](#)

    single topology [RC-280](#)

customizing routes [RC-303](#)

default routes [RC-273](#)

description [RC-265](#)

enabling [RC-277](#)

enabling multicast-intact [RC-308](#)

functional overview [RC-267](#)

grouped configuration [RC-268](#)

IP fast reroute [RC-313](#)

IPv6 routing [RC-269](#)

Level 1 or Level 2 routing, configuration [RC-277](#)

LSP flooding

    controlling [RC-285](#)

    lifetime maximum [RC-270](#)

    limiting [RC-269](#)

MPLS LDP IS-IS synchronization [RC-306](#)

MPLS TE

    configuring [RC-295](#)

- description [RC-273](#)
- multi-instance IS-IS [RC-272](#)
- multitopology, configuring [RC-284](#)
- nonstop forwarding [RC-272](#)
  - configuring [RC-289](#)
- overload bit
  - configuring [RC-270](#)
  - on router [RC-273](#)
- policy attach points
  - default originate [RC-399](#)
  - inter-area-propagate [RC-399](#)
  - redistribute [RC-399](#)
- priority for prefixes added to RIB [RC-311](#)
- restrictions, configuring [RC-266](#)
- set SPF interval [RC-301](#)
- single topology
  - configuring [RC-280](#)
- single-topology
  - IPv6 support [RC-271](#)
  - tagging IS-IS interface routes [RC-309](#)
- IS-IS address family submode [RC-268](#)
- ispf command [RC-302](#)
- is-type command [RC-279](#)

---

## K

- keychain command [RC-115](#)
- keychains [RC-39](#)

---

## L

- L2TPv3 tunnels [RC-40](#)
- label-allocation-mode per-ce command [RC-106](#)
- link-state advertisement (LSA)
  - OSPFv2 [RC-181](#)
  - OSPFv3 [RC-182, RC-190](#)
- log adjacency changes command [RC-196, RC-251, RC-299](#)
- log-neighbor-changes command [RC-164](#)

- log-neighbor-warnings command [RC-164](#)
- loop-free alternate [RC-276, RC-317, RC-323, RC-331](#)
- lsp-check-interval command [RC-286](#)
- LSP flooding
  - controlling [RC-285](#)
  - lifetime maximum [RC-270](#)
  - limiting [RC-269](#)
  - mesh group configuration [RC-270](#)
    - on specific interfaces [RC-270](#)
- lsp-gen-interval command [RC-286](#)
- lsp-interval command [RC-287](#)
- lsp-mtu command [RC-286](#)
- lsp-password command [RC-292, RC-294](#)
- lsp-refresh-interval command [RC-286](#)

---

## M

- maximum paths command [RC-154](#)
- maximum-paths command [RC-110, RC-305](#)
- maximum prefix command [RC-97](#)
- maximum-prefix command [RC-155](#)
- max-lsp-lifetime command [RC-286](#)
- mesh-group command [RC-287](#)
- Message Digest 5 (MD5) authentication [RC-193](#)
- message-digest-key command [RC-205](#)
- message statistics [RC-322](#)
- metric-style wide command [RC-297, RC-310, RC-312](#)
- metro-zero-accept command [RC-341](#)
- MPLS LDP IGP synchronization [RC-275](#)
- MPLS LDP IS-IS synchronization [RC-306](#)
- mpls ldp sync command [RC-307](#)
- MPLS TE (Multiprotocol Label Switching traffic engineering) configuring
  - IS-IS [RC-295](#)
  - OSPFv2 [RC-226](#)
- MPLS TE forwarding adjacency [RC-276](#)
- MPLS TE interarea tunnels [RC-276](#)
- mpls traffic-eng area command [RC-228](#)
- mpls traffic-eng command [RC-296](#)

- mpls traffic-engineering multicast-intact command [RC-309](#)
- mpls traffic-eng router-id command [RC-227](#), [RC-296](#)
- MPLS VPN carrier supporting carrier [RC-38](#)
- multi-area adjacency
  - configuring [RC-244](#)
  - interface attributes and limitations [RC-191](#)
  - overview [RC-191](#)
- multi-area-interface command [RC-245](#)
- multicast-intact [RC-274](#)
  - OSPFv2 [RC-190](#)
- multicast topology [RC-275](#)
- multicast VPN [RC-41](#)
- multi-instance IS-IS [RC-272](#)
- multiprotocol BGP [RC-32](#)
- multitopology
  - configuring [RC-284](#)
  - example [RC-315](#)

## N

- NBMA networks [RC-179](#)
- neighbor address family command [RC-8](#), [RC-269](#)
- neighbor command [RC-8](#), [RC-268](#), [RC-338](#)
- neighbor command (OSPFv2, OSPFv3) [RC-202](#)
- neighbor-group command [RC-76](#)
- neighbors
  - adjacency (OSPFv2) [RC-181](#)
  - maximum limits (BGP) [RC-4](#)
- net command [RC-278](#)
- network command [RC-59](#), [RC-201](#)
- next-hop-self command [RC-84](#)
- nonstop forwarding, configuring (OSPFv2) [RC-224](#)
- not-so-stubby area [RC-178](#)
- nsf command [RC-289](#), [RC-340](#)
- nsf interface-expires command [RC-290](#)
- nsf interface-timer command [RC-290](#)
- nsf interval command [RC-226](#)
- nsf lifetime command [RC-290](#)

- nssa command [RC-198](#)

## O

- ospf area configuration submode
  - dead-interval command [RC-201](#)
  - default-cost command [RC-198](#)
  - hello interval command [RC-201](#)
  - interface command [RC-196](#)
  - network command [RC-201](#)
  - nssa command [RC-198](#)
  - range command [RC-214](#)
  - stub command [RC-198](#)
- ospf area submode
  - authentication message-digest command [RC-211](#)
  - virtual-link command [RC-211](#)
- ospf interface configuration submode
  - log adjacency changes [RC-196](#), [RC-251](#)
  - neighbor command [RC-202](#)
- OSPFv2 (Open Shortest Path First Version 2)
  - authentication, configuring [RC-204](#)
  - Cisco IOS XR OSPFv3 and OSPFv2 differences [RC-175](#)
  - CLI (command-line interface) inheritance [RC-176](#)
  - configuration
    - MPLS TE [RC-226](#)
    - neighbors, nonbroadcast networks [RC-199](#)
    - sham-links [RC-233](#)
  - configuration and operation, verifying [RC-254](#)
  - default route [RC-181](#)
  - description [RC-171](#)
  - Designate Router (DR) [RC-181](#)
  - enabling [RC-194](#)
  - functional overview [RC-173](#)
  - instance and router ID [RC-179](#)
  - LSA
    - controlling the frequency [RC-207](#)
    - on an OSPF ABR [RC-213](#)
    - types [RC-181](#)



- MD5 authentication [RC-180](#)
- MPLS TE, configuring [RC-226](#)
- neighbors, adjacency [RC-181](#)
- neighbors, nonbroadcast networks, configuring [RC-199](#)
- nonstop forwarding
  - configuring [RC-224](#)
  - description [RC-187](#)
- policy attach points
  - area-in [RC-396](#)
  - area-out [RC-396](#)
  - default originate [RC-395](#)
  - redistribute [RC-395, RC-397](#)
- route authentication methods
  - key rollover [RC-180](#)
  - MD5 [RC-180](#)
  - plain text [RC-180](#)
  - strategies [RC-180](#)
- route redistribution
  - configuring [RC-215](#)
  - description [RC-186](#)
- sham-link [RC-233](#)
- Shortest Path First (SPF) throttling
  - configuring [RC-219](#)
  - description [RC-186](#)
- stub and not-so-stubby area types, configuring [RC-197](#)
- supported OSPF network types
  - NBMA networks [RC-179](#)
  - point to point networks [RC-179](#)
- virtual link
  - creating [RC-209](#)
  - transit area [RC-183](#)
- OSPFv2 (Open Shortest Path First version 2)
  - enabling multicast-intact [RC-237](#)
- OSPFv3 (Open Shortest Path First Version 3) [RC-219](#)
  - Cisco IOS XR OSPFv3 and OSPFv2 differences [RC-175](#)
  - CLI inheritance [RC-176](#)
  - configuration
    - neighbors, nonbroadcast networks [RC-199](#)
    - SPF throttling [RC-219](#)
  - configuration and operation, verifying [RC-254](#)
  - default route [RC-181](#)
  - description [RC-171](#)
  - enabling [RC-194](#)
  - functional overview [RC-173](#)
  - graceful restart [RC-230](#)
  - instance and router ID [RC-179](#)
  - load balancing [RC-191](#)
  - LSA
    - controlling frequency [RC-207](#)
    - on an OSPF ABR [RC-213](#)
    - types [RC-182](#)
  - neighbors
    - nonbroadcast networks, configuring [RC-199](#)
  - policy attach points
    - default originate [RC-397](#)
    - redistribute [RC-395, RC-397](#)
  - routes, redistribute [RC-215](#)
  - SPF (Shortest Path First) throttling configuring [RC-219](#)
  - stub and not-so-stubby area types, configuring [RC-197](#)
  - virtual link, description [RC-183](#)
- ospfv3 area configuration submode
  - dead-interval command [RC-201](#)
  - default-cost command [RC-198](#)
  - hello interval command [RC-201](#)
  - interface command [RC-196](#)
  - network command [RC-201](#)
  - nssa command [RC-198](#)
  - range command [RC-214](#)
  - stub command [RC-198](#)
- OSPFv3 Graceful Restart feature [RC-188](#)
  - adjacency [RC-189](#)
  - displaying information [RC-233](#)
- ospfv3 interface configuration submode
  - log adjacency changes [RC-196, RC-251](#)



neighbor command [RC-202](#)  
 output-delay command [RC-340](#)  
 overload bit  
   configuration [RC-270](#)  
   on router [RC-273](#)

---

## P

passive-interface command [RC-343](#)  
 password command [RC-102](#), [RC-107](#)  
 PCE extensions to OSPFv2 [RC-193](#)  
 PIM [RC-274](#)  
 point-to-point networks [RC-179](#)  
 poison-reverse command [RC-341](#)  
 policy, modifying  
   attached [RC-406](#)  
   nonattached [RC-406](#)

---

## R

range command [RC-214](#)  
 rd command [RC-100](#)  
 receive version command [RC-339](#)  
 redistribute command [RC-65](#), [RC-113](#), [RC-154](#), [RC-160](#), [RC-162](#),  
[RC-217](#)  
 redistribute isis command [RC-304](#)  
 redistribute maximum-prefix command [RC-154](#)  
 retain route-target command [RC-105](#)  
 retransmit-interval command [RC-287](#)  
 retransmit-throttle-interval command [RC-287](#)  
 RFC 1142, OSI IS-IS Intra-domain Routing  
 Protocol [RC-318](#)  
 RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and  
 Dual Environments [RC-318](#)  
 RFC 1587, Not So Stubby Area (NSSA) [RC-263](#)  
 RFC 1700, Assigned Numbers [RC-134](#)  
 RFC 1771 [RC-417](#)  
 RFC 1793, OSPF over demand circuit [RC-263](#)  
 RFC 1997, BGP Communities Attribute [RC-134](#)

RFC 2328, OSPF Version 2 [RC-175](#), [RC-263](#)  
 RFC 2370, LSA [RC-263](#)  
 RFC 2385, Protection of BGP Sessions via the TCP MD5  
 Signature Option [RC-134](#)  
 RFC 2439, BGP Route Flap Damping [RC-134](#)  
 RFC 2453, RIP Version 2 [RC-334](#), [RC-351](#)  
 RFC 2545, Use of BGP-4 Multiprotocol Extensions for  
 IPv6 Inter-Domain Routing [RC-134](#)  
 RFC 2740, OSPFv3 [RC-263](#)  
 RFC 2740 OSPFv3 [RC-175](#)  
 RFC 2763, Dynamic Hostname Exchange Mechanism for  
 IS-IS [RC-318](#)  
 RFC 2796, BGP Route Reflection - An Alternative to Full  
 Mesh IBGP [RC-134](#)  
 RFC 2858, Multiprotocol Extensions for BGP-4 [RC-134](#)  
 RFC 2918, Route Refresh Capability for BGP-4 [RC-134](#)  
 RFC 2966, Domain-wide Prefix Distribution with  
 Two-Level IS-IS [RC-318](#)  
 RFC 2973, IS-IS Mesh Groups [RC-318](#)  
 RFC 3065, Autonomous System Confederations for  
 BGP [RC-134](#)  
 RFC 3137, Router Advertisement [RC-263](#)  
 RFC 3277, IS-IS Transient Blackhole Avoidance [RC-318](#)  
 RFC 3373, Three-Way Handshake for IS-IS Point-to-Point  
 Adjacencies [RC-318](#)  
 RFC 3392, Capabilities Advertisement with  
 BGP-4 [RC-134](#)  
 RFC 3509 [RC-263](#)  
 RFC 3567, IS-IS Cryptographic Authentication [RC-318](#)  
 RFC 3623, OSPFv3 [RC-263](#)  
 RFC 4136, OSPF [RC-263](#)  
 RFC 4271 [RC-134](#)  
 RFC 4360 [RC-417](#)  
 RFC 4364, BGP/MPLS IP Virtual Private Networks  
 (VPNs) [RC-134](#)  
 RFC 4576, L3VPN [RC-263](#)  
 RFC 4577, Sham Link [RC-263](#)  
 RFC 4724, Graceful Restart Mechanisms for BGP [RC-134](#)  
 RFC 4750, OSPF Version 2 [RC-263](#)  
 RFC 4811, OSPF [RC-263](#)  
 RFC 4812, OSPF [RC-263](#)

- RFC 4813, OSPF [RC-263](#)
- RIB (Routing Information Base)
  - administrative distance [RC-321](#)
  - data structures in BGP and other protocols [RC-321](#)
  - deploying [RC-324](#)
  - description [RC-319](#)
  - examples [RC-327](#)
  - functional overview [RC-320](#)
  - IPv4 and IPv6 support [RC-322](#)
  - monitoring [RC-324](#)
  - prerequisites [RC-320](#)
- RIB quarantining [RC-323](#)
- RIB statistics [RC-322](#)
- RIP (Routing Information Protocol)
  - policy attach points
    - global-inbound [RC-404](#)
- RIP (Routing Information Protocol)
  - administrative distance [RC-336](#)
  - bandwidth [RC-342](#)
  - benefits [RC-335](#)
  - CIDR [RC-334](#)
  - control or prevent routing updates [RC-342](#)
  - filter network updates [RC-342](#)
  - hop count [RC-334](#)
  - metrics [RC-334](#)
  - policy attach points
    - default-information originate [RC-404](#)
    - global-inbound [RC-404](#)
    - interface-inbound [RC-405](#)
    - interface-outbound [RC-405](#)
    - redistribute [RC-404](#)
  - redistribution [RC-336](#)
  - route policy creation [RC-344](#)
  - route timers [RC-335](#)
  - routing loops [RC-342](#)
  - routing policy options [RC-337](#)
  - split horizon, enabling IP [RC-335](#)
  - VLSMs (variable-length subnet masks) [RC-334](#)
  - WAN link [RC-342](#)
- RIP v2 supported features [RC-334](#)
- route dampening [RC-33](#)
- route-policy (BGP) command [RC-47](#), [RC-76](#), [RC-81](#)
- route-policy command [RC-45](#), [RC-80](#), [RC-103](#), [RC-108](#), [RC-117](#), [RC-152](#), [RC-156](#), [RC-344](#), [RC-345](#), [RC-409](#)
- route-policy configuration submode
  - set eigrp-metric command [RC-156](#)
  - set rip-metric command [RC-345](#)
- route-policy pass-all command [RC-20](#)
- route policy submode [RC-45](#)
  - end-policy command [RC-45](#)
  - See* route-policy command
- router bgp command [RC-8](#), [RC-268](#)
- router bgp neighbor group address family configuration mode
  - address family command [RC-10](#)
- route redistribution (OSPFv2, OSPFv3) [RC-186](#)
- route-reflector-client command [RC-79](#)
- route reflectors [RC-34](#)
- router eigrp address family configuration submode [RC-158](#)
  - default-metric command [RC-150](#)
  - distance command [RC-150](#)
  - log-neighbor-changes command [RC-164](#)
  - log-neighbor-warnings command [RC-164](#)
  - maximum paths command [RC-154](#)
  - maximum-prefix command [RC-155](#)
  - redistribute command [RC-154](#), [RC-160](#), [RC-162](#)
  - redistribute maximum-prefix command [RC-154](#)
  - timers nsf route-hold command [RC-154](#)
- router eigrp command [RC-149](#)
- router eigrp configuration mode
  - route-policy command [RC-152](#)
- router eigrp configuration submode
  - address-family ipv4 command [RC-149](#)
  - interface command [RC-150](#)
- router eigrp interface configuration submode
  - bandwidth-percent command [RC-151](#)
  - holdtime command [RC-150](#)
  - summary-address command [RC-153](#)

- router-id command [RC-195](#)
- router identifier [RC-3](#)
- router isis address family submode
  - default-information originate command [RC-304](#)
  - distance command [RC-305](#)
  - ispf command [RC-302](#)
  - maximum-paths command [RC-305](#)
  - metric-style wide command [RC-297, RC-310, RC-312](#)
  - mpls ldp sync command [RC-307](#)
  - mpls traffic-eng command [RC-296](#)
  - mpls traffic-eng multicast-intact command [RC-309](#)
  - mpls traffic-eng router-id command [RC-296](#)
  - redistribute isis command [RC-304](#)
  - set-attached-bit command [RC-306](#)
  - single-topology command [RC-283](#)
  - spf-interval command [RC-302](#)
  - spf prefix-priority command [RC-312](#)
  - summary-prefix command [RC-305](#)
- router isis configuration submode
  - address-family command [RC-283](#)
  - ignore-lsp-errors command [RC-287](#)
  - is-type command [RC-279](#)
  - log adjacency changes command [RC-299](#)
  - lsp-check-interval command [RC-286](#)
  - lsp-gen-interval command [RC-286](#)
  - lsp-mtu command [RC-286](#)
  - lsp-password command [RC-292, RC-294](#)
  - lsp-refresh-interval command [RC-286](#)
  - max-lsp-lifetime command [RC-286](#)
  - net command [RC-278](#)
  - nsf command [RC-289](#)
  - nsf interface-expires command [RC-290](#)
  - nsf interface-timer command [RC-290](#)
  - nsf lifetime command [RC-290](#)
  - set-overload-bit command [RC-304](#)
- router isis interface address family submode
  - tag command [RC-310](#)
- router isis interface configuration submode
  - address-family command [RC-283](#)
  - circuit-type command [RC-283, RC-313, RC-314](#)
  - csnp-interval command [RC-287](#)
  - hello-interval command [RC-299](#)
  - hello-multiplier command [RC-299](#)
  - hello-padding command [RC-299](#)
  - hello-password command [RC-293, RC-295, RC-300](#)
  - ipv4 address command [RC-282](#)
  - ipv6 address command [RC-282](#)
  - ipv6 enable command [RC-282](#)
  - lsp-interval command [RC-287](#)
  - mesh-group command [RC-287](#)
  - retransmit-interval command [RC-287](#)
  - retransmit-throttle-interval command [RC-287](#)
- router isis interface submode
  - mesh-group command [RC-287](#)
- router ospf command [RC-195](#)
- router ospf configuration submode [RC-195](#)
  - area command [RC-196](#)
  - authentication command [RC-205](#)
  - message-digest-key command [RC-205](#)
  - mpls traffic-eng area command [RC-228](#)
  - mpls traffic-eng router-id command [RC-227](#)
  - nsf interval command [RC-226](#)
  - redistribute command [RC-217](#)
  - router-id command [RC-195](#)
  - summary-prefix command [RC-218](#)
  - timers lsa gen-interval command [RC-208](#)
  - timers lsa group-pacing command [RC-209](#)
  - timers lsa min-interval command [RC-208](#)
  - timers throttle spf command [RC-220, RC-221, RC-222, RC-252](#)
- router ospfv3 command [RC-195](#)
- router ospfv3 configuration submode
  - area command [RC-196](#)
  - redistribute command [RC-217](#)
  - router-id command [RC-195](#)
  - summary-prefix command [RC-218](#)
  - timers lsa gen-interval command [RC-208](#)
  - timers lsa group-pacing command [RC-209](#)

- timers lsa min-interval command [RC-208](#)
- timers throttle spf command [RC-220](#), [RC-221](#), [RC-222](#), [RC-252](#)
- router rib command [RC-327](#)
- router rib configuration submode
  - address-family command [RC-327](#)
- router rip command [RC-338](#)
- router rip configuration submode
  - auto-summary command [RC-340](#)
  - broadcast-for-v2 command [RC-338](#)
  - interface command [RC-339](#)
  - neighbor command [RC-338](#)
  - nsf command [RC-340](#)
  - output-delay command [RC-340](#)
  - timers basic command [RC-340](#)
- router rip interface configuration submode
  - metro-zero-accept command [RC-341](#)
  - passive-interface command [RC-343](#)
  - poison-reverse command [RC-341](#)
  - receive version command [RC-339](#)
  - route-policy command [RC-344](#)
  - send version command [RC-339](#)
  - split-horizon disable command [RC-341](#)
- router static command [RC-429](#)
- routes
  - customizing (IS-IS) [RC-303](#)
  - default
    - IS-IS [RC-273](#)
    - OSPFv2 [RC-181](#)
  - redistribute (OSPFv2, OSPFv3) [RC-215](#)
  - redistribute IS-IS routes example [RC-315](#)
- route tags [RC-274](#)
- routing components
  - Area Border Routers (ABRs) [RC-178](#)
  - Autonomous System Boundary Routers (ASBRs) [RC-178](#)
  - autonomous systems [RC-177](#)
  - backbone area [RC-177](#)
  - Designated Router (DR) [RC-181](#)
  - interior routers [RC-179](#)
  - not-so-stubby area [RC-178](#)
  - stub area [RC-178](#)
- routing domain confederation [RC-34](#)
- routing policy [RC-20](#)
  - attaching to BGP neighbor [RC-410](#)
  - configuration elements, editing [RC-407](#)
  - defining [RC-409](#)
  - defining (example) [RC-414](#)
  - enforcing, BGP [RC-20](#)
  - implementing
    - prerequisites [RC-354](#)
  - inbound (example) [RC-414](#)
  - modifying [RC-412](#)
  - modular inbound (example) [RC-415](#)
  - statements
    - action [RC-372](#)
    - disposition [RC-370](#)
    - elseif [RC-372](#)
    - if [RC-372](#)
    - remark [RC-370](#)
- RPL (routing policy language)
  - Boolean operators, types [RC-373](#)
  - components [RC-360](#)
  - overview [RC-355](#)
  - policy
    - attributes
    - modification [RC-366](#)
    - parameterization [RC-364](#)
      - Boolean operator precedence [RC-366](#)
      - configuration basics [RC-363](#)
      - default drop disposition [RC-367](#)
      - definitions [RC-363](#)
      - statement processing [RC-368](#)
      - statements, types [RC-370](#)
      - verification [RC-368](#)
  - structure
    - as-path-set, inline set form [RC-357](#)
    - as-path-set, named set form [RC-357](#)

community-set, inline set form [RC-357](#)  
 community-set, named set form [RC-357](#)  
 extended community set, inline form [RC-358](#)  
 extended community set, named form [RC-358](#)  
 names [RC-355](#)  
 prefix-set [RC-359](#)  
 sets [RC-356](#)

## S

security ttl command [RC-252](#)  
 send-community-ebgp command [RC-86](#)  
 send version command [RC-339](#)  
 set-attached-bit command [RC-306](#)  
 set eigrp-metric command [RC-156](#)  
 set-overload-bit command [RC-304](#)  
 set rip-metric command [RC-345](#)  
 sham-links [RC-233](#)  
 show bgp af-group command [RC-16](#), [RC-17](#)  
 show bgp cidr-only command [RC-124](#), [RC-125](#)  
 show bgp community command [RC-124](#), [RC-125](#), [RC-126](#)  
 show bgp count-only command [RC-125](#)  
 show bgp flap-statistics command [RC-68](#)  
 show bgp flap-statistics reexp command [RC-68](#)  
 show bgp flap-statistics route-policy command [RC-69](#)  
 show bgp inheritance command [RC-16](#)  
 show bgp neighbor command [RC-15](#)  
 show bgp neighbor-group command [RC-18](#), [RC-125](#)  
 show bgp neighbors command [RC-124](#)  
 show bgp paths command [RC-124](#)  
 show bgp reexp command [RC-124](#)  
 show bgp session-group command [RC-17](#)  
 show bgp summary command [RC-125](#)  
 show eigrp accounting command [RC-165](#)  
 show eigrp interfaces command [RC-165](#)  
 show eigrp neighbors command [RC-159](#), [RC-166](#)  
 show eigrp topology command [RC-166](#)  
 show eigrp traffic command [RC-166](#)  
 show ip route connected command [RC-329](#)  
 show isis command [RC-279](#)  
 show isis database command [RC-288](#)  
 show isis database-log command [RC-288](#)  
 show isis interface command [RC-301](#)  
 show isis lsp-log command [RC-288](#)  
 show isis mpls command [RC-297](#)  
 show isis mpls traffic-eng adjacency-log command [RC-297](#)  
 show isis mpls traffic-eng advertisements command [RC-297](#)  
 show isis neighbors command [RC-301](#)  
 show isis spf-log command [RC-303](#)  
 show isis topology command [RC-284](#)  
 show ospf command [RC-211](#)  
 show ospfv3 command [RC-211](#)  
 show protocols eigrp command [RC-166](#)  
 show rpl route-policy command [RC-413](#)  
 show running-config command [RC-290](#)  
 shutdown command [RC-102](#), [RC-119](#)  
 single-topology  
   command [RC-283](#)  
   configuring example [RC-315](#)  
   IPv6 support [RC-271](#)  
   set SPF interval [RC-301](#)  
 site-of-origin command [RC-108](#)  
 soft-reconfiguration inbound always command [RC-93](#)  
 spf-interval command [RC-302](#)  
 spf prefix-priority command [RC-312](#)  
 SPF throttling, configuring  
   OSPFv2 (Open Shortest Path First Version 2) [RC-219](#)  
 split-horizon disable command [RC-341](#)  
 static router submode  
   address-family command [RC-424](#), [RC-426](#), [RC-428](#),  
   [RC-431](#)  
   maximum path command [RC-430](#)  
   vrf command [RC-424](#), [RC-426](#), [RC-428](#), [RC-431](#)  
 static routes  
   administrative distance [RC-421](#)  
   associating with VRF [RC-430](#)  
   connected [RC-421](#)

defined [RC-420](#)  
 floating [RC-423, RC-425, RC-427](#)  
 maximum number of routes [RC-429](#)  
 MPLS layer 3 VPN [RC-421](#)  
 recursive [RC-422](#)  
 specified [RC-422](#)  
 stub area [RC-178](#)  
 stub area types, configuring (OSPFv3) [RC-197](#)  
 stub command [RC-158, RC-198](#)  
 summary-address command [RC-153](#)  
 summary-prefix command [RC-218, RC-305](#)  
 synchronization  
     MPLS LDP IGP [RC-275](#)

---

## T

table-policy command [RC-72](#)  
 tag command [RC-310](#)  
 task group prerequisite [RC-334](#)  
 timers basic command [RC-340](#)  
 timers bgp command [RC-50](#)  
 timers command [RC-102](#)  
 timers lsa gen-interval command [RC-208](#)  
 timers lsa group-pacing command [RC-209](#)  
 timers lsa min-interval command [RC-208](#)  
 timers nsf route-hold command [RC-154](#)  
 timers throttle spf command [RC-220, RC-221, RC-222, RC-252](#)

---

## U

update groups  
     BGP configuration [RC-22](#)  
     BGP update generation [RC-22](#)  
     monitor [RC-126](#)  
 update-source command [RC-102](#)  
 use command [RC-76](#)

---

## V

virtual link  
     transit area (OSPFv2) [RC-183](#)  
 virtual-link command [RC-211](#)  
 VPNv4 address family command [RC-9](#)  
 VPNv4 and VPNv6 over the IP core using L2TPv3 tunnels [RC-40](#)  
 VPNv6 address family command [RC-9](#)  
 VRF address family command [RC-8](#)  
 vrf address family submode  
     export route-target command [RC-98](#)  
 VRF command [RC-8](#)  
 VRF neighbor address family command [RC-9](#)  
 VRF neighbor command [RC-8](#)

---

## W

weight command [RC-56](#)