cisco.



Routing Configuration Guide for Cisco NCS 5000 Series Routers, IOS XR Release 7.2.x

First Published: 2020-08-01

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 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.

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned 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. (1721R)

© 2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

P R E F A C E	Preface xi
	Changes to This Document xi
	Changes to This Document xi
	Communications, Services, and Additional Information xi
CHAPTER 1	New and Changed Routing Features 1
	New and Changed Routing Features 1
	New and Changed Routing Features 1
CHAPTER 2	Implementing IS-IS 3
	Enable IS-IS and Configure Level 1 or Level 2 Routing 3
	Customize Routes for IS-IS 5
	Set Priority for Adding Prefixes to RIB 9
	IS-IS Interfaces 10
	Tag IS-IS Interface Routes 10
	Limit LSP Flooding 13
	Control LSP Flooding for IS-IS 13
	IS-IS Authentication 17
	Configure Authentication for IS-IS 18
	Configure Keychains for IS-IS 19
	ISIS NSR 21
	Configuring ISIS-NSR 21
	Configuring IS-IS Adjacency Stagger 23
	IS-IS Overload Bit Avoidance 23
	Configure IS-IS Overload Bit Avoidance 24
	References for IS-IS 25

IS-IS Functional Overview 25 Default Routes 25 Overload Bit on Router 26 Overload Bit Configuration During Multitopology Operation 26 Attached Bit on an IS-IS Instance 26 IS-IS Support for Route Tags 26 Flood Blocking on Specific Interfaces 27 Maximum LSP Lifetime and Refresh Interval 27 Mesh Group Configuration 27 Multi-Instance IS-IS 27

CHAPTER 3 Implementing OSPF 29

Prerequisites for Implementing OSPF 30
Enable OSPF 30
Verify OSPF Configuration and Operation 33
Stub Area 35
Not-so-Stubby Area 35
Configure Stub and Not-So-Stubby Area Types 35
Neighbors and Adjacency for OSPF 38
Configure Neighbors for Nonbroadcast Networks 39
Authentication Strategies 43
Configure Authentication at Different Hierarchical Levels for OSPF Version 2 43
Control Frequency That Same LSA Is Originated or Accepted for OSPF 46
Virtual Link and Transit Area for OSPF 48
Create Virtual Link 49
Summarize Subnetwork LSAs on OSPF ABR 54
Route Redistribution for OSPF 56
Redistribute Routes into OSPF 56
OSPF Shortest Path First Throttling 59
Configure OSPF Shortest Path First Throttling 60
Graceful Restart for OSPFv3 62
Configure OSPFv3 Graceful Restart 63
Display Information About Graceful Restart 64
OSPFv2OSPF SPF Prefix Prioritization 65

Configure OSPFv2 OSPF SPF Prefix Prioritization **66** Multi-Area Adjacency for OSPF Version 2 70 Configure Multi-area Adjacency 70 Label Distribution Protocol IGP Auto-configuration for OSPF 72 Configure Label Distribution Protocol IGP Auto-configuration for OSPF 73 Configure LDP IGP Synchronization: OSPF 74 IGP link state **76** OSPF Authentication Message Digest Management **76** Configure Authentication Message Digest Management for OSPF 77 OSPF Link-State Database Overload Protection 80 References for OSPF 81 OSPF Functional Overview 81 Comparison of Cisco IOS XR Software OSPFv3 and OSPFv2 82 OSPF Hierarchical CLI and CLI Inheritance 83 OSPF Routing Components 83 Autonomous Systems 84 Areas 84 Routers 85 OSPF Process and Router ID 85 Supported OSPF Network Types 86 Route Authentication Methods for OSPF 86 Plain Text Authentication 86 MD5 Authentication 86 Key Rollover 87 OSPF FIB Download Notification 87 Designated Router (DR) for OSPF 87 Default Route for OSPF 87 Link-State Advertisement Types for OSPF Version 2 87 Link-State Advertisement Types for OSPFv3 88 Passive Interface 89 Modes of Graceful Restart Operation 90 Restart Mode 90 Helper Mode 90 Protocol Shutdown Mode 91

Load Balancing in OSPF Version 2 and OSPFv3 92 Path Computation Element for OSPFv2 92 Management Information Base (MIB) for OSPFv3 92 OSPFv3 Timers Update 93

CHAPTER 4 Implementing and Monitoring RIB 95

Verify RIB Configuration Using Routing Table 95 Verify Networking and Routing Problems 96 Disable RIB Next-hop Dampening 98 Enable RCC and LCC On-demand Scan 99 Enable RCC and LCC Background Scan 100 References for RIB 102 RIB Data Structures in BGP and Other Protocols 102 RIB Administrative Distance 103 RIB Statistics 103 RIB Quarantining 104 Route and Label Consistency Checker 104

CHAPTER 5 Implementing RIP 107

Prerequisites for Implementing RIP 107 Information About Implementing RIP 107 RIP Functional Overview 107 Split Horizon for RIP 108 Route Timers for RIP 109 Route Redistribution for RIP 109 Default Administrative Distances for RIP 110 Routing Policy Options for RIP 110 Authentication Using Keychain in RIP **111** In-bound RIP Traffic on an Interface **112** Out-bound RIP Traffic on an Interface 112 How to Implement RIP **112** Enabling RIP 113 Customizing RIP 114 Control Routing Information 116

Creating a Route Policy for RIP Configuring RIP Authentication Keychain Configuring RIP Authentication Keychain for IPv4 Interface on a Non-default VRF Configuring RIP Authentication Keychain for IPv4 Interface on Default VRF Configuration Examples for Implementing RIP Configuring a Basic RIP Configuration: Example Configuring RIP on the Provider Edge: Example Adjusting RIP Timers for each VRF Instance: Example Configuring Redistribution for RIP: Example Configuring Passive Interfaces and Explicit Neighbors for RIP: Example

CHAPTER 6 Implementing Routing Policy 127

Restrictions for Implementing Routing Policy 127 Define Route Policy 128 Attach Routing Policy to BGP Neighbor 130 Modify Routing Policy Using Text Editor 131 References for Routing Policy 134 Routing Policy Language 134 Routing Policy Language Overview 134 Routing Policy Language Structure 135 Routing Policy Language Components 142 Routing Policy Language Usage 142 Policy Definitions 144 Parameterization 145 Parameterization at Attach Points 146 Global Parameterization 147 Semantics of Policy Application 147 Boolean Operator Precedence 147 Multiple Modifications of Same Attribute 148 When Attributes Are Modified 149 Default Drop Disposition 149 Control Flow 149 Policy Verification 150

Policy Statements 152 Remark 152 Disposition 152 Action 154 If 155 Boolean Conditions 156 apply 157 Attach Points 157 BGP Policy Attach Points 158 **OSPF** Policy Attach Points 175 OSPFv3 Policy Attach Points 179 IS-IS Policy Attach Points 180 Nondestructive Editing of Routing Policy 181 Attached Policy Modification 181 Nonattached Policy Modification 182 Editing Routing Policy Configuration Elements 182 Hierarchical Policy Conditions 184 Apply Condition Policies 184 Nested Wildcard Apply Policy 187 Match Aggregated Route 188 Remove Private AS in Inbound Policy 188 **Implementing Static Routes** 189 Restrictions for Implementing Static Routes 189 Configure Static Route 189 Floating Static Routes 191 Configure Floating Static Route 191 Change Maximum Number of Allowable Static Routes 193 Default VRF 194 References for Static Routes 195 Static Route Functional Overview 195 Default Administrative Distance 195

Directly Connected Routes 196

Floating Static Routes 196

CHAPTER 7

Fully Specified Static Routes 196 **Recursive Static Routes** 196 Dynamic ECMP 197 IPv4 Multicast Static Routes 197 Configure Multicast Static Routes 198 **CHAPTER 8 Implementing Integrated Routing and Bridging** 201 Bridge-Group Virtual Interface 201 Sub-Interfaces with BVI interface in Bridge Domain 202 Supported Features on a BVI 202 BVI Interface and Line Protocol States 202 Prerequisites for Configuring IRB 203 Restrictions for Configuring IRB 203 How to Configure IRB 203 Configuring the Bridge Group Virtual Interface 204 Configuration Guidelines 204 Configuring the Layer 2 AC Interfaces 206 Configuring a Bridge Group and Assigning Interfaces to a Bridge Domain 207 Configure the Layer 2 AC Sub-Interfaces 209 Configure a Bridge Group and Assign Sub-Interfaces to a Bridge Domain 209 Associating the BVI as the Routed Interface on a Bridge Domain 210 Displaying Information About a BVI 211 Additional Information on IRB 212 Packet Flows Using IRB 212 Packet Flows When Host A Sends to Host B on the Bridge Domain 213 Packet Flows When Host A Sends to Host C From the Bridge Domain to a Routed Interface 213 Packet Flows When Host C Sends to Host B From a Routed Interface to the Bridge Domain 213 Configuration Examples for IRB 214 Basic IRB Configuration: Example 214 IPv4 Addressing on a BVI Supporting Multiple IP Networks: Example 214

CHAPTER 9 Implementing BFD 215

BFD over Bundle **215** Configure BFD over Bundle **216** BFD over Bundle and BFD over Logical Bundle 218
Enabling BFD Sessions on Bundle Members 218
Specifying the BFD Destination Address on a Bundle 219
Configuring the Minimum Thresholds for Maintaining an Active Bundle 220
Configuring BFD Packet Transmission Intervals and Failure Detection Times on a Bundle 222
Configure BFD over Bundles IETF Mode Support on a Per Bundle Basis 223



Preface



Note This product has reached end-of-life status. For more information, see the End-of-Life and End-of-Sale Notices.

The Routing Configuration Guide for Cisco NCS 5000 Series Routers preface contains these sections:

- Changes to This Document, on page xi
- Changes to This Document, on page xi
- · Communications, Services, and Additional Information, on page xi

Changes to This Document

This table lists the technical changes made to this document since it was first released.

Table 1: Changes to This Document

Date	Summary
January 2020	Initial release of this document.

Changes to This Document

This table lists the technical changes made to this document since it was first released.

Table 2: Changes to This Document

Date	Summary
August 2020	Initial release of this document.

Communications, Services, and Additional Information

• To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.

- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPIEK

New and Changed Routing Features

This table summarizes the new and changed feature information for the *Routing Configuration Guide for Cisco NCS 5000 Series Routers*, and tells you where they are documented.

- New and Changed Routing Features, on page 1
- New and Changed Routing Features, on page 1

New and Changed Routing Features

Table 3: Routing Features Added or Modified in IOS XR Release 7.1.x

Feature	Description	Changed in Release	Where Documented
None	No new features introduced	Not applicable	Not applicable

New and Changed Routing Features

Table 4: Routing Features Added or Modified in IOS XR Release 7.2.x

Feature	Description	Changed in Release	Where Documented
None	No new features introduced	Not applicable	Not applicable



Implementing IS-IS

Integrated Intermediate System-to-Intermediate System (IS-IS), Internet Protocol Version 4 (IPv4), is a standards-based Interior Gateway Protocol (IGP). The Cisco software implements the IP routing capabilities described in International Organization for Standardization (ISO)/International Engineering Consortium (IEC) 10589 and RFC 1195, and adds the standard extensions for single topology and multitopology IS-IS for IP Version 6 (IPv6).

This module describes how to implement IS-IS (IPv4 and IPv6) on your Cisco IOS XR network.

- Enable IS-IS and Configure Level 1 or Level 2 Routing, on page 3
- Customize Routes for IS-IS, on page 5
- Set Priority for Adding Prefixes to RIB, on page 9
- IS-IS Interfaces, on page 10
- Limit LSP Flooding, on page 13
- IS-IS Authentication, on page 17
- ISIS NSR, on page 21
- Configuring IS-IS Adjacency Stagger, on page 23
- IS-IS Overload Bit Avoidance, on page 23
- References for IS-IS, on page 25

Enable IS-IS and Configure 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.



Note Users can configure the **no max-metric** command only with levels 1 or 2, that is, **no max-metric level {1|2}** in order to view the result in the output of the **show configuration** command. Else, the maximum metric configuration is not displayed in the output. This behavior is observed before committing the configuration to the router.

Before you begin

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. Use the commit or end command.
- 6. show isis [instance instance-id] protocol

DETAILED STEPS

 Step 1
 configure

 Example:
 RP/0/RP0/CPU0:router# configure

 Enters mode.
 Enters mode.

 Step 2
 router isis instance-id

 Example:
 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.

 • 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 is-type router configuration command.

Step 3 net *network-entity-title*

Example:

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.

- Specify a NET for each routing instance if you are configuring multi-instance IS-IS.
- This example configures a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110.00.
- 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.
- Step 4 is-type { level-1 | level-1-2 | level-2-only }

Example:

```
RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only
```

(Optional) Configures the system type (area or backbone router).

- By default, every IS-IS instance acts as a level-1-2 router.
- The **level-1** 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 **level-1-2** router in the area.
- The **level-2-only** 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 **level-1-2** routers.
- The level-1-2 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.
- Step 5 Use the commit or end command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

 Step 6
 show isis [instance instance-id] protocol

Example:

RP/0/RP0/CPU0:router# show isis protocol

(Optional) Displays summary information about the IS-IS instance.

Customize 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

- 1. configure
- **2.** router isis *instance-id*
- **3**. set-overload-bit [on-startup { *delay* | wait-for-bgp }] [level { 1 | 2 }]
- 4. address-family { ipv4 | ipv6 } [unicast]
- 5. default-information originate [route-policy route-policy-name]
- 6. redistribute isis *instance* [level-1 | level-2 | level-1-2] [metric *metric*] [metric-type { internal | external }] [policy *policy-name*]
- **7.** Do one of the following:

```
• summary-prefix address / prefix-length [level { 1 | 2 }]
```

• 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.** Use the **commit** or **end** command.

DETAILED STEPS

Step 1	configure		
	Example:		
	RP/0/RP0/CPU0:router# configure		
	Enters mode.		
Step 2	router isis instance-id		
	Example:		
	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.		
	• 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 is-type command.		
Step 3	set-overload-bit [on-startup { delay wait-for-bgp }] [level { 1 2 }]		
	Example:		
	RP/0/RP0/CPU0:router(config-isis)# set-overload-bit		
	(Optional) Sets the overload bit.		
	Note 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	address-family { ipv4 ipv6 } [unicast]		
	Example:		
	RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast		
	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.		
Step 5	default-information originate [route-policy route-policy-name]		
	Example:		
	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.		
	• The route-policy keyword and <i>route-policy-name</i> argument specify the conditions under which the IPv4 or IPv6 default route is advertised.		

- If the **route-policy** keyword is omitted, then the IPv4 or IPv6 default route is unconditionally advertised at Level 2.
- Step 6
 redistribute isis instance [level-1 | level-2 | level-1-2][metric metric][metric-type { internal | external }][policy policy-name]

Example:

RP/0/RP0/CPU0:router(config-isis-af)# redistribute isis 2 level-1

(Optional) Redistributes routes from one IS-IS instance into another instance.

- In this example, an IS-IS instance redistributes Level 1 routes from another IS-IS instance.
- **Step 7** Do one of the following:
 - summary-prefix address / prefix-length [level {1 | 2}]
 - summary-prefix ipv6-prefix | prefix-length [level { 1 | 2 }]

Example:

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:xxxx::/24 level 1

(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.

• This example specifies an IPv4 address and mask.

or

- 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.
- 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.

```
Step 8 maximum-paths route-number
```

Example:

RP/0/RP0/CPU0:router(config-isis-af)# maximum-paths 16

(Optional) Configures the maximum number of parallel paths allowed in a routing table.

Step 9 distance weight [address / prefix-length [route-list-name]]

Example:

RP/0/RP0/CPU0:router(config-isis-af)# distance 90

(Optional) Defines the administrative distance assigned to routes discovered by the IS-IS protocol.

• A different administrative distance may be applied for IPv4 and IPv6.

Step 10 set-attached-bit

Example:

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 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

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 HundredGigE 0/3/0/0
  address-family ipv4 unicast
 1
 1
 router isis 2
  is-type level-1
  net 49.0002.0001.0001.0002.00
  address-family ipv4 unicast
   set
-attached-bit
 interface HundredGigE 0/1/0/0
  address-family ipv4 unicast
```

Set Priority for Adding Prefixes to 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]
- 4. metric-style wide [transition][level {1 | 2}]
- **5.** spf prefix-priority [level {1 | 2}] { critical | high | medium } { access-list-name | tag tag }
- **6.** Use the **commit** or **end** command.

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router isis instance-id
	Example:
	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	address-family { ipv4 ipv6 } [unicast]
	Example:
	RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 4	metric-style wide [transition] [level { 1 2 }]
	Example:
	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	<pre>spf prefix-priority [level { 1 2 }] { critical high medium } { access-list-name tag tag } Example:</pre>
	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 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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.

Tag 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]
- 4. metric-style wide [transition] [level $\{1 \mid 2\}$]
- 5. exit
- **6. interface** *type number*
- 7. address-family { ipv4 | ipv6 } [unicast]
- 8. tag tag
- 9. Use the commit or end command.
- **10.** show isis [ipv4 | ipv6 | afi-all] [unicast | safi-all] route [detail]

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router isis instance-id
	Example:
	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	address-family { ipv4 ipv6 } [unicast]
	Example:
	RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 4	metric-style wide [transition] [level { 1 2 }]
	Example:
	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	exit
	Example:
	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	interface type number
	Example:
	RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/1/0/3
	Enters interface configuration mode.
Step 7	address-family { ipv4 ipv6 } [unicast]
	Example:
	<pre>RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast</pre>
	Specifies the IPv4 or IPv6 address family, and enters address family configuration mode.

Step 8 tag tag Example: RP/0/RP0/CPU0:router(config-isis-if-af) # tag 3 Sets the value of the tag to associate with the advertised connected route. Step 9 Use the commit or end command. commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: · Yes — Saves configuration changes and exits the configuration session. · No —Exits the configuration session without committing the configuration changes. · Cancel —Remains in the configuration session, without committing the configuration changes. Step 10 show isis [ipv4 + ipv6 + afi-all] [unicast + safi-all] route [detail]

RP/0/RP0/CPU0:router# show isis ipv4 route detail

Displays tag information. Verify that all tags are present in the RIB.

Tagging Routes: Example

The following example shows how to tag routes.

```
route-policy isis-tag-55
end-policy
1
route-policy isis-tag-555
 if destination in (5.5.5.0/24 \text{ eq } 24) then
    set tag 555
    pass
  else
    drop
  endif
end-policy
router static
address-family ipv4 unicast
 0.0.0/0 2.6.0.1
  5.5.5.0/24 Null0
 !
1
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

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.

Control 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. **Isp-check-interval** seconds [level {1 | 2}]
- 5. lsp-gen-interval { [initial-wait initial | secondary-wait secondary | maximum-wait maximum] ... } [level { 1 | 2 }]
- **6. Isp-mtu** bytes [level { 1 | 2 }]
- 7. max-lsp-lifetime seconds [level $\{1 \mid 2\}$]
- 8. ignore-lsp-errors disable
- **9. interface** *type interface-path-id*
- **10.** Isp-interval milliseconds [level $\{1 \mid 2\}$]
- **11.** csnp-interval seconds [level $\{1 \mid 2\}$]

- **12.** retransmit-interval seconds [level { 1 | 2 }]
- **13.** retransmit-throttle-interval *milliseconds* [level { 1 | 2 }]
- **14.** mesh-group { number | blocked }
- **15.** Use the **commit** or **end** command.
- **16.** show isis interface [type interface-path-id | 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 $\{1 \mid 2\}$]
- **19.** show isis database-log [level $\{1 \mid 2\}$]

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router isis instance-id
	Example:
	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.
	• You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.
Step 3	lsp-refresh-interval seconds [level { 1 2 }]
	Example:
	RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 10800
	(Optional) Sets the time between regeneration of LSPs that contain different sequence numbers
	• The refresh interval should always be set lower than the max-lsp-lifetime command.
Step 4	lsp-check-interval seconds [level { 1 2 }]
	Example:
	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.
	• This operation is costly in terms of CPU and so should be configured to occur infrequently.
Step 5	lsp-gen-interval { [initial-wait initial secondary-wait secondary maximum-wait maximum] } [level { 1 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5 secondary-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.

- 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.
- **Step 6 Isp-mtu** bytes [level $\{1 \mid 2\}$]

Example:

RP/0/RP0/CPU0:router(config-isis)# lsp-mtu 1300

(Optional) Sets the maximum transmission unit (MTU) size of LSPs.

Step 7 max-lsp-lifetime seconds [level { 1 | 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis) # max-lsp-lifetime 11000

(Optional) Sets the initial lifetime given to an LSP originated by the router.

• This is the amount of time that the LSP persists in the database of a neighbor unless the LSP is regenerated or refreshed.

Step 8 ignore-lsp-errors disable

Example:

RP/0/RP0/CPU0:router(config-isis) # ignore-lsp-errors disable

(Optional) Sets the router to purge LSPs received with checksum errors.

Step 9 *interface type interface-path-id*

Example:

RP/0/RP0/CPU0:router(config-isis) # interface HundredGigE 0/1/0/3

Enters interface configuration mode.

Step 10Isp-intervalmilliseconds[level { 1 | 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis-if) # lsp-interval 100

(Optional) Configures the amount of time between each LSP sent on an interface.

Step 11csnp-intervalseconds[level { 1 | 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis-if)# csnp-interval 30 level 1

(Optional) Configures the interval at which periodic CSNP packets are sent on broadcast interfaces.

- Sending more frequent CSNPs means that adjacent routers must work harder to receive them.
- Sending less frequent CSNP means that differences in the adjacent routers may persist longer.

Step 12 retransmit-interval seconds [level { 1 | 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis-if)# retransmit-interval 60

(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.

RP/0/RP0/CPU0:router(config-isis-if)# retransmit-interval 60

Step 13 retransmit-throttle-interval milliseconds [level { 1 | 2 }]

Example:

RP/0/RP0/CPU0:router(config-isis-if)# retransmit-throttle-interval 1000

(Optional) Configures the amount of time between retransmissions on each LSP on a point-to-point interface.

• This time is usually greater than or equal to the **lsp-interval** 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.

Step 14 mesh-group { number | blocked }

Example:

RP/0/RP0/CPU0:router(config-isis-if) # mesh-group blocked

(Optional) Optimizes LSP flooding in NBMA networks with highly meshed, point-to-point topologies.

This command is appropriate only for an NBMA network with highly meshed, point-to-point topologies.

Step 15 Use the commit or end command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- **Cancel** —Remains in the configuration session, without committing the configuration changes.
- Step 16
 show isis interface [type interface-path-id | level { 1 | 2 }] [brief]

 Example:

RP/0/RP0/CPU0:router# show isis interface HundredGigE 0/1/0/1 brief

(Optional) Displays information about the IS-IS interface.

Step 17	show isis [instance instance-id] database [level { 1 2 }] [detail summary verbose] [* lsp-id]
	Example:
	RP/0/RP0/CPU0:router# show isis database level 1
	(Optional) Displays the IS-IS LSP database.
Step 18	show isis [instance instance-id] lsp-log [level {1 2}]
	Example:
	RP/0/RP0/CPU0:router# show isis lsp-log
	(Optional) Displays LSP log information.
Step 19	show isis database-log [level $\{1 \mid 2\}$]
	Example:
	RP/0/RP0/CPU0:router# show isis database-log level 1
	(Optional) Display IS-IS database log information.

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. ey 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).

Configure 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 interface-path-id*
- 5. hello-password { hmac-md5 | text } { clear | encrypted } password [level { 1 | 2 }] [send-only]
- 6. Use the commit or end command.

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure Enters mode.

Step 2 router isis instance-id

Example:

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.

• You can change the level of routing to be performed by a particular routing instance by using the is-type command.

$\label{eq:step3} Step 3 \qquad lsp-password \ \{\ hmac-md5 \ | \ text \ \} \ \{\ clear \ | \ encrypted \ \} \ password \ [\ level \ \ \{\ 1 \ | \ 2 \ \}] \ [\ send-only \] \ send-only \] \ [\ send-only \] \ [\ send-only \] \ send-only \] \ [\ send-only \] \ [\ send-only \] \ send-only \] \ [\ send-only \] \ send-only \ send-only \] \ send-only \] \ send-only \ send-only \] \ send-only \ send-only \ send-only \] \ send-only \ send-only \] \ send-only \ send-only \] \ send-only \ se$

Example:

RP/0/RP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1

Configures the LSP authentication password.

- The **hmac-md5** keyword specifies that the password is used in HMAC-MD5 authentication.
- The text keyword specifies that the password uses cleartext password authentication.
- The clear keyword specifies that the password is unencrypted when entered.
- The encrypted keyword specifies that the password is encrypted using a two-way algorithm when entered.

- The level 1 keyword sets a password for authentication in the area (in Level 1 LSPs and Level SNPs).
- The level 2 keywords set a password for authentication in the backbone (the Level 2 area).
- The **send-only** keyword adds authentication to LSP and sequence number protocol data units (SNPs) when they are sent. It does not authenticate received LSPs or SNPs.
- The **snp send-only** keyword adds authentication to SNPs when they are sent. It does not authenticate received SNPs.
- **Note** To disable SNP password checking, the **snp send-only** keywords must be specified in the **lsp-password** command.
- **Step 4** interface type interface-path-id

Example:

RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/3

Enters interface configuration mode.

 Step 5
 hello-password { hmac-md5 | text } { clear | encrypted } password [level { 1 | 2 }] [send-only]

 Example:

```
RP/0/RP0/CPU0:router(config-isis-if)#hello-password text clear mypassword
```

Configures the authentication password for an IS-IS interface.

Step 6 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

- end —Prompts user to take one of these actions:
 - Yes Saves configuration changes and exits the configuration session.
 - No —Exits the configuration session without committing the configuration changes.
 - Cancel Remains in the configuration session, without committing the configuration changes.

Configure 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

1. configure

- 2. router isis instance-id
- **3.** I sp-password keychain keychain-name [level { 1 | 2 }] [send-only] [snp send-only]
- **4. interface** *type interface-path-id*
- 5. hello-password keychain keychain-name [level {1 | 2}] [send-only]
- 6. Use the commit or end command.

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router isis instance-id
	Example:
	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.
	• You can change the level of routing to be performed by a particular routing instance by using the is-type command.
Step 3	l sp-password keychain keychain-name [level { 1 2 }] [send-only] [snp send-only]
	Example:
	RP/0/RP0/CPU0:router(config-isis)# lsp-password keychain isis_a level 1
	Configures the keychain.
Step 4	interface type interface-path-id
	Example:
	RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/1/0/3
	Enters interface configuration mode.
Step 5	hello-password keychain keychain-name [level { 1 2 }] [send-only]
	Example:
	RP/0/RP0/CPU0:router(config-isis-if)#hello-password keychain isis_b
	Configures the authentication password for an IS-IS interface.
Step 6	Use the commit or end command.
	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.

- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

ISIS NSR

Non Stop Routing (NSR) suppresses IS-IS routing changes for devices with redundant route processors during processor switchover events (RP failover or ISSU), reducing network instability and downtime. When Non Stop Routing is used, switching from the active to standby RP have no impact on the other IS-IS routers in the network. All information needed to continue the routing protocol peering state is transferred to the standby processor prior to the switchover, so it can continue immediately upon a switchover.

To preserve routing across process restarts, NSF must be configured in addition to NSR.

Configuring ISIS-NSR

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router isis instance-id
	Example:
	RP/0/RP0/CPU0:router(config)# router isis 1
	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.
Step 3	nsr
	Example:
	RP/0/RP0/CPU0:router(config-isis)# nsr
	Configures the NSR feature.
Step 4	Use the commit or end command.
	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.
	• No —Exits the configuration session without committing the configuration changes.
	• Cancel — Remains in the configuration session, without committing the configuration changes.

Step 5 show isis nsr adjacency

Example:

RP/0/RP0/CPU0:router
show isis nsr adjacency
System Id Interface SNPA State Hold Changed NSF IPv4 BFD IPv6 BFD
R1-v1S Nii0 *PtoP* Up 83 00:00:33 Yes None None

Displays adjacency information.

Step 6 show isis nsr status

Example:

RP/0/RP0/CPU0:router										
router#show isis nsr status										
IS-IS test NSR(v1a) STATUS	Ready):									
	V1	Standby	V2 Active	V2 Standby						
SYNC STATUS:		TRUE	FALSE(0)	FALSE(0)						
PEER CHG COUNT:		1	0	0						
UP TIME:	0 (D:03:12	not up	not up						

Displays the NSR status information.

Step 7 show isis nsr statistics

Example:

RP/0/RP0/CPU0:router router#show isis nsr statistics									
IS-IS test NSR(v1a) MANDATORY STATS :									
10 10 0000 101((114) 11	V1 Active	V1 Standby	V2 Active	V2					
Standby									
L1 ADJ:	0	0	0						
0									
L2 ADJ:	2	2	0						
0									
LIVE INTERFACE:	4	4	0						
0									
PTP INTERFACE:	1	1	0						
0									
LAN INTERFACE:	2	2	0						
	1	1	0						
LOOPBACK INTERFACE:	1	Ţ	0						
TE Tunnel:	1	1	0						
0	Ť	Ŧ	0						
TE LINK:	2	2	0						
0	_	-	Ũ						
NSR OPTIONAL STATS :									
L1 LSP:	0	0	0						
0									
L2 LSP:	4	4	0						
0									
IPV4 ROUTES:	3	3	0						
0									
IPV6 ROUTES:	4	4	0						
0									

Shows number of ISIS adjacencies, lsps, routes, tunnels, Te links on active and standby routers.

Configuring IS-IS Adjacency Stagger

Certain events like process restart or reload can involve a significant processing overhead. Updating routing tables with all adjacencies, maintaining them, and synchronizing the database with each adjacent router requires a lot of bandwidth. These processes may require large number of packets being sent and/or received, depending on the state of the database on the routers. If packets are dropped in any direction, it can lead to an unstable state.

We cannot prevent events like process restart or reload, but we can handle such events better by limiting the number of adjacencies that area being established simultaneously. To limit the number of adjacencies from getting established simultaneously, you can configure adjacency stagger. By configuring IS-IS adjacency stagger, you can specify the initial number neighbourhood routers from which adjacencies can fully form after a process restart or reload. If you configure IS-IS adjacency stagger, you can also specify the subsequent number of simultaneous neighbors that are allowed to form adjacency.

Restrictions

- IS-IS adjacency stagger is only supported on point-to-point interfaces and not on LAN interfaces.
- IS-IS adjacency stagger is not supported with NSF (non-stop forwarding) mechanisms.

Configuration Example

To configure IS-IS adjacency stagger on a point-to-point interface, you must use the following configuration steps:

- 1. Configure IS-IS.
- 2. Configure adjacency stagger.

Router(config-isis) # commit

Configuration

```
/* Enter the global configuration mode and configure IS-IS */
Router# config
Router(config)# router isis 1
/* Configure IS-IS adjacency stagger */
Router(config-isis)# adjacency stagger 2 3
```

IS-IS Overload Bit Avoidance

The IS-IS overload bit avoidance feature allows network administrators to prevent label switched paths (LSPs) from being disabled when a router in that path has its Intermediate System-to-Intermediate System (IS-IS) overload bit set.

When the IS-IS overload bit avoidance feature is activated, all nodes with the overload bit set, including head nodes, mid nodes, and tail nodes, are ignored, which means that they are still available for use with label switched paths (LSPs).



Note

The IS-IS overload bit avoidance feature does *not* change the default behavior on nodes that have their overload bit set if those nodes are not included in the path calculation (PCALC).

The IS-IS overload bit avoidance feature is activated using the following command:

mpls traffic-eng path-selection ignore overload

The IS-IS overload bit avoidance feature is deactivated using the no form of this command:

no mpls traffic-eng path-selection ignore overload

When the IS-IS overload bit avoidance feature is deactivated, nodes with the overload bit set cannot be used as nodes of last resort.

Configure IS-IS Overload Bit Avoidance

This task describes how to activate IS-IS overload bit avoidance.

Before you begin

The IS-IS overload bit avoidance feature is valid only on networks that support the following features:

- MPLS
- IS-IS

SUMMARY STEPS

- 1. configure
- 2. mpls traffic-eng path-selection ignore overload

DETAILED STEPS

Step 1 configure Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 2 mpls traffic-eng path-selection ignore overload Example:

RP/0/RP0/CPU0:router(config) # mpls traffic-eng path-selection ignore overload

Activates IS-IS overload bit avoidance.

Configuring IS-IS Overload Bit Avoidance: Example

The following example shows how to activate IS-IS overload bit avoidance:

```
config mpls traffic-eng path-selection ignore overload
```

The following example shows how to deactivate IS-IS overload bit avoidance:

```
config
  no mpls traffic-eng path-selection ignore overload
```

References for IS-IS

This section provides additional conceptual information on IS-IS. It includes the following topics:

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).

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.

Multiple IS-IS instances can exist on the same physical interface. However, you must configure different instance-id for every instance that shares the same physical interface.

Alternatively, you can also create dot1q sub-interfaces and configure each dot1q sub-interface to different IS-IS instances.

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 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.

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.

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.

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).

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.

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.

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.

Multi-Instance IS-IS

You can configure up to 16 IS-IS instances. MPLS can run on multiple IS-IS processes as long as the processes run on different sets of interfaces. Each interface may be associated with 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 cause 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.

I



Implementing OSPF

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 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 software. The term "OSPF" implies both versions of the routing protocol, unless otherwise noted.

Note GTSM TTL Security is not supported.

- Prerequisites for Implementing OSPF, on page 30
- Enable OSPF, on page 30
- Verify OSPF Configuration and Operation, on page 33
- Stub Area, on page 35
- Neighbors and Adjacency for OSPF, on page 38
- Authentication Strategies, on page 43
- Control Frequency That Same LSA Is Originated or Accepted for OSPF, on page 46
- Virtual Link and Transit Area for OSPF, on page 48
- Summarize Subnetwork LSAs on OSPF ABR, on page 54
- Route Redistribution for OSPF, on page 56
- OSPF Shortest Path First Throttling, on page 59
- Graceful Restart for OSPFv3, on page 62
- OSPFv2OSPF SPF Prefix Prioritization, on page 65
- Multi-Area Adjacency for OSPF Version 2, on page 70
- Label Distribution Protocol IGP Auto-configuration for OSPF, on page 72
- IGP link state, on page 76
- OSPF Authentication Message Digest Management, on page 76
- OSPF Link-State Database Overload Protection, on page 80
- References for OSPF, on page 81

Prerequisites for Implementing OSPF

The following are prerequisites for implementing OSPF:

- Configuration tasks for OSPFv3 assume that you are familiar with IPv6 addressing and basic configuration. See the *Implementing Network Stack IPv4 and IPv6* in the *Cisco 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.

Enable 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.

Before you begin

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.** Do one of the following:
 - router ospf process-name
 - router ospfv3 process-name
- **3.** router-id { router-id }
- 4. area area-id
- **5.** interface type interface-path-id
- 6. Repeat Step 5 for each interface that uses OSPF.
- 7. log adjacency changes [detail] [enable | disable]
- 8. Use the commit or end command.

DETAILED STEPS

Step 1 configure

	Example:	
	RP/0/RP0/CPU0:router# configure	
	Enters mode.	
Step 2	Do one of the following:	
	 router ospf process-name router ospfv3 process-name 	
	Example:	
	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.	
	Note The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.	
Step 3	router-id { router-id }	
	Example:	
	RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3	
	Configures a router ID for the OSPF process.	
	Note We recommend using a stable IP address as the router ID.	
Step 4	area area-id	
	Example:	
	RP/0/RP0/CPU0:router(config-ospf)# area 0	
	Enters area configuration mode and configures an area for the OSPF process.	
	• Backbone areas have an area ID of 0.	
	Nonbackbone areas have a nonzero area ID.	
	• The <i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.23. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.	2.

Step 5 interface type interface-path-id

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/1/0/3

Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.

- **Step 6** Repeat Step 5 for each interface that uses OSPF.
- Step 7 log adjacency changes [detail][enable | disable] Example:

RP/0/RP0/CPU0:router(config-ospf-ar-if)# log adjacency changes detail

(Optional) Requests notification of neighbor changes.

- By default, this feature is enabled.
- The messages generated by neighbor changes are considered notifications, which are categorized as severity Level 5 in the **logging console** command. The **logging console** command controls which severity level of messages are sent to the console. By default, all severity level messages are sent.
- **Step 8** Use the **commit** or **end** command.

commit — Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

Enable OSPF: Example

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.

```
interface TenGigE 0/3/0/0
 ip address 10.1.2.1 255.255.255.0
negotiation auto
!
interface TenGigE 0/3/0/1
ip address 10.1.3.1 255.255.255.0
negotiation auto
L
router ospf 1
router-id 10.2.3.4
area O
interface TenGigE 0/3/0/0
T.
area 1
interface TenGigE 0/3/0/1
!
1
```

Verify 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 interface-path-id
- 5. **show** { **ospf** | **ospfv3** } [*process-name*] [**vrf** *vrf-name*] [*area-id*] **interface** [*type interface-path-id*]
- 6. show { ospf | ospfv3 }[process-name] [area-id] neighbor [t ype interface- path-id] [neighbor-id] [detail]
- 7. clear { ospf | ospfv3 }[process-name] process
- 8. clear {ospf|ospfv3[process-name] redistribution
- 9. clear {ospf|ospfv3[process-name] routes
- **10.** clear {ospf|ospfv3[*process-name*] vrf [*vrf-name*|all] {process |redistribution|routes|statistics [interface type interface-path-id|message-queue|neighbor]}
- **11.** clear { ospf | ospfv3 }[process-name] statistics [neighbor [type interface-path-id] [ip-address]]

DETAILED STEPS

Step 1	<pre>show { ospf ospfv3 } [process-name]</pre>
	Example:
	RP/0/RP0/CPU0:router# show ospf group1
	(Optional) Displays general information about OSPF routing processes.
Step 2	<pre>show { ospf ospfv3 } [process-name] border-routers [router-id]</pre>
	Example:
	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	<pre>show { ospf ospfv3 } [process-name] database</pre>
	Example:
	RP/0/RP0/CPU0:router# show ospf group2 database
	(Optional) Displays the lists of information related to the OSPF database for a specific router.
	• The various forms of this command deliver information about different OSPF LSAs.
Step 4	show { ospf ospfv3 } [<i>process-name</i>] [<i>area-id</i>] flood-list interface <i>type interface-path-id</i>

Example:

RP/0/RP0/CPU0:router# show ospf 100 flood-list interface TenGigE 0/3/0/0

(Optional) Displays a list of OSPF LSAs waiting to be flooded over an interface.

 Step 5
 show { ospf | ospfv3 } [process-name] [vrf vrf-name] [area-id] interface [type interface-path-id]

 Example:

RP/0/RP0/CPU0:router# show ospf 100 interface TenGigE 0/3/0/0
(Optional) Displays OSPF interface information.

 Step 6
 show { ospf | ospfv3 }[process-name] [area-id] neighbor [t ype interface- path-id] [neighbor-id] [detail]

 Example:

RP/0/RP0/CPU0:router# show ospf 100 neighbor

(Optional) Displays OSPF neighbor information on an individual interface basis.

Step 7 clear { ospf | ospfv3 }[process-name] process

Example:

RP/0/ /CPU0:router# clear ospf 100 process

(Optional) Resets an OSPF router process without stopping and restarting it.

Step 8 clear {ospf|ospfv3[process-name] redistribution

Example:

RP/0/RP0/CPU0:router#clear ospf 100 redistribution

Clears OSPF route redistribution.

Step 9 clear {ospf|ospfv3[process-name] routes

Example:

RP/0/RP0/CPU0:router#clear ospf 100 routes

Clears OSPF route table.

 Step 10
 clear {ospf|ospfv3[process-name] vrf [vrf-name|all] {process |redistribution|routes|statistics [interface type interface-path-id|message-queue|neighbor]}

Example:

RP/0/RP0/CPU0:router#clear ospf 100 vrf vrf_1 process

Clears OSPF route table.

 Step 11
 clear { ospf | ospfv3 }[process-name] statistics [neighbor [type interface-path-id] [ip-address]]

 Example:

```
RP/0/RP0/CPU0:router# clear ospf 100 statistics
```

(Optional) Clears the OSPF statistics of neighbor state transitions.

Stub Area

A stub area is an area that does not accept route advertisements 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.

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. Use the **capability type7 translate zero-forward-addr** command to enable the translation of Type-7 LSA with a zero forwarding address into Type-5 LSA on the NSSA ABR router. This command is supported only for OSPFv3. 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.

Configure 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.** Do one of the following:
 - router ospf process-name
 - router ospfv3 process-name
- 3. capability type7 translate zero-forward-addr
- 4. router-id { router-id }
- 5. area area-id
- **6.** Do one of the following:
 - stub [no-summary]
 - nssa [no-redistribution][default-information-originate][no-summary]

- **7.** Do one of the following:
 - stub
 - nssa
- 8. default-cost cost
- 9. Use the commit or end command.
- **10.** Repeat this task on all other routers in the stub area or NSSA.

DETAILED STEPS

Step 1	configure	2
	Example:	
	RP/0/RP0	/CPU0:router# configure
	Enters mo	ode.
Step 2	Do one of	f the following:
		er ospf process-name er ospfv3 process-name
	Example:	
	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 C	OSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.
	Note	The process-name argument is any alphanumeric string no longer than 40 characters.
Step 3	capabilit	y type7 translate zero-forward-addr
	Example:	
		/CPU0:router(config)# router ospfv3 1 /CPU0:router(config-ospfv3)# capability type7 translate zero-forward-addr
	Enables t	he translation of Type-7 LSA with a zero forwarding address into Type-5 LSA on the NSSA ABR router.
	Note	You can configure capability type7 translate zero-forward-addr command only in router ospfv3 configuration mode. This command is not supported for the OSPF process.
Step 4	router-id	{ router-id }
	Example:	

RP/0/RP0/CPU0:router(config-ospf) # router-id 192.168.4.3

	C C	
	-	s a router ID for the OSPF process.
	Note	We recommend using a stable IP address as the router ID.
Step 5	area area	ı-id
	Example:	
	RP/0/RP0/	CPU0:router(config-ospf)# area 1
	Enters area	a configuration mode and configures a nonbackbone area for the OSPF process.
		<i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address on.
Step 6	Do one of	the following:
		[no-summary] [no-redistribution] [default-information-originate] [no-summary]
	Example:	
	RP/0/RP0/	CPU0:router(config-ospf-ar)# stub no summary
	or	
	RP/0/RP0/	CPU0:router(config-ospf-ar)# nssa no-redistribution
	Defines the	e nonbackbone area as a stub area.
	-	fy the no-summary keyword to further reduce the number of LSAs sent into a stub area. This keyword nts the ABR from sending summary link-state advertisements (Type 3) in the stub area.
	or	
	Defines an	area as an NSSA.
Step 7	Do one of	the following:
	• stub	
	• nssa	
	Example:	
	RP/0/RP0/	CPU0:router(config-ospf-ar)# stub
	or	
	RP/0/RP0/	CPU0:router(config-ospf-ar)# nssa
	(Optional)	Turns off the options configured for stub and NSSA areas.
	defau	a configured the stub and NSSA areas using the optional keywords (no-summary , no-redistribution , ilt-information-originate , and no-summary) in Step 5, you must now reissue the stub and nssa commands ut the keywords—rather than using the no form of the command.

• For example, the **no nssa default-information-originate** form of the command changes the NSSA area into a normal area that inadvertently brings down the existing adjacencies in that area.

Step 8 default-cost cost

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)#default-cost 15

(Optional) Specifies a cost for the default summary route sent into a stub area or an NSSA.

- Use this command only on ABRs attached to the NSSA. Do not use it on any other routers in the area.
- The default cost is 1.

Step 9 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

- end —Prompts user to take one of these actions:
 - Yes Saves configuration changes and exits the configuration session.
 - No —Exits the configuration session without committing the configuration changes.
 - Cancel —Remains in the configuration session, without committing the configuration changes.
- **Step 10** Repeat this task on all other routers in the stub area or NSSA.

Configuring a Stub area: 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 TenGigE 0/2/0/1
area 1
stub
interface TenGigE 0/2/0/0
```

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.

Configure Neighbors for Nonbroadcast Networks

This task explains how to configure neighbors for a nonbroadcast network. This task is optional.

Before you begin

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

- 1. configure
- **2.** Do one of the following:
 - router ospf process-name
 - router ospfv3 process-name
- **3.** router-id { router-id }
- 4. area area-id
- **5**. network { broadcast | non-broadcast }
- 6. dead-interval seconds
- 7. hello-interval seconds
- **8.** interface type interface-path-id
- **9.** Do one of the following:
 - neighbor *ip-address* [priority *number*] [poll-interval seconds][cost number]
 - 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 interface-path-id
- **13.** Do one of the following:
 - neighbor *ip-address* [priority *number*] [poll-interval *seconds*][cost *number*] [database-filter [all]]
 - 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.** Use the **commit** or **end** command.

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 2	Do one or	f the following:
	• rout	er ospf process-name
	• rout	er ospfv3 process-name
	Example:	
	RP/0/RP0	/CPU0:router(config)# router ospf 1
	or	
	RP/0/RP0	/CPU0:router(config)# router ospfv3 1
	Enables (DSPF routing for the specified routing process and places the router in router configuration mode.
	or	
	Enables (DSPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.
	Note	The process-name argument is any alphanumeric string no longer than 40 characters.
Step 3	router-id	{ router-id }
	Example:	
	RP/0/RP0	/CPU0:router(config-ospf)# router-id 192.168.4.3
	Configure	es a router ID for the OSPF process.
	Note	We recommend using a stable IP address as the router ID.
Step 4	area area	a-id
•	Example:	
	RP/0/RP0	/CPU0:router(config-ospf)# area 0
		ea configuration mode and configures an area for the OSPF process.
		example configures a backbone area.
	• The	<i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address
Step 5	network	{ broadcast non-broadcast }

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# network non-broadcast

Configures the OSPF network type to a type other than the default for a given medium.

• The example sets the network type to NBMA.

Step 6 dead-interval seconds

Example:

L

	RP/0/RP0/CPU0:router(config-ospf-ar)# dead-interval 40
	(Optional) Sets the time to wait for a hello packet from a neighbor before declaring the neighbor down.
Step 7	hello-interval seconds
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# hello-interval 10
	(Optional) Specifies the interval between hello packets that OSPF sends on the interface.
Step 8	interface type interface-path-id
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/2/0/0
	Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.
	• 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.
Step 9	Do one of the following:
	 neighbor ip-address [priority number] [poll-interval seconds] [cost number] neighbor ipv6-link-local-address [priority number] [poll-interval seconds] [cost number] [database-filter [all]]
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.20.20.1 priority 3 poll-interval 15
	or
	<pre>RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor fe80::3203:a0ff:fe9d:f3fe</pre>
	Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.

or

Configures the link-local IPv6 address of OSPFv3 neighbors.

- The *ipv6-link-local-address* argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.
- The **priority** 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.
- Neighbors with no specific cost configured assumes the cost of the interface, based on the cost command.
- The **database-filter** keyword filters outgoing LSAs to an OSPF neighbor. If you specify the **all** keyword, incoming and outgoing LSAs are filtered.
- **Step 10** Repeat Step 9 for all neighbors on the interface.

Step 11 exit

Example:

RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit

Enters area configuration mode.

Step 12 interface type interface-path-id

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/3/0/0

Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.

• 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.

Step 13 Do one of the following:

neighbor *ip-address* [priority *number*] [poll-interval *seconds*] [cost *number*] [database-filter [all]]
neighbor *ipv6-link-local-address* [priority *number*] [poll-interval *seconds*] [cost *number*] [database-filter

```
[ all ]]
```

Example:

```
RP/0/
/CPU0:router(config-ospf-ar)# neighbor 10.34.16.6
```

or

RP/0/
/CPU0:router(config-ospf-ar)# neighbor fe80::3203:a0ff:fe9d:f3f

Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.

or

Configures the link-local IPv6 address of OSPFv3 neighbors.

- The *ipv6-link-local-address* argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.
- The **priority** 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.
- Neighbors with no specific cost configured assumes the cost of the interface, based on the cost command.
- The **database-filter** keyword filters outgoing LSAs to an OSPF neighbor. If you specify the **all** 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 traffic disruption or routing loops.
- **Step 14** Repeat Step 13 for all neighbors on the interface.
- **Step 15** Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

• Yes — Saves configuration changes and exits the configuration session.

- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

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.

Configure 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.

Before you begin

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 OSPF Hierarchical CLI and CLI Inheritance, on page 83 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 | LINE}
- 6. area area-id
- 7. interface type interface-path-id
- **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 interface-path-id
- **13.** Repeat Step 12 for each interface that must communicate, using the same authentication.
- **14.** interface type interface-path-id
- **15.** authentication [message-digest | null]
- **16.** Use the **commit** or **end** command.

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router ospf process-name
	Example:
	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.
	Note The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.
Step 3	<pre>router-id { router-id }</pre>
	Example:
	RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3
	Configures a router ID for the OSPF process.
Step 4	authentication [message-digest null]
	Example:
	RP/0/RP0/CPU0:router(config-ospf)#authentication message-digest
	Enables MD5 authentication for the OSPF process.
	• This authentication type applies to the entire router process unless overridden by a lower hierarchical level such as the area or interface.
Step 5	message-digest-key key-id md5 { key clear key encrypted key LINE}
	Example:
	RP/0/RP0/CPU0:router(config-ospf)#message-digest-key 4 md5 yourkey
	Specifies the MD5 authentication key for the OSPF process.
	• The neighbor routers must have the same key identifier.
Step 6	area area-id

I

	Example:
	RP/0/RP0/CPU0:router(config-ospf)# area 0
	Enters area configuration mode and configures a backbone area for the OSPF process.
Step 7	interface type interface-path-id
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/1/0/3
	Enters interface configuration mode and associates one or more interfaces to the backbone area.
	• All interfaces inherit the authentication parameter values specified for the OSPF process (Step 4, Step 5, and Step 6).
Step 8	Repeat Step 7 for each interface that must communicate, using the same authentication.
Step 9	 exit
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# exit
	Enters area OSPF configuration mode.
Step 10	area area-id
	Example:
	RP/0/RP0/CPU0:router(config-ospf)# area 1
	Enters area configuration mode and configures a nonbackbone area 1 for the OSPF process.
	• 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 11	authentication [message-digest null]
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# authentication
	Enables Type 1 (plain text) authentication that provides no security.
	• The example specifies plain text authentication (by not specifying a keyword). Use the authentication-key command in interface configuration mode to specify the plain text password.
Step 12	interface type interface-path-id
	Example:
	RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/1/0/0
	Enters interface configuration mode and associates one or more interfaces to the nonbackbone area 1 specified in Step 7.

• All interfaces configured inherit the authentication parameter values configured for area 1.

Step 13 Repeat Step 12 for each interface that must communicate, using the same authentication.

Step 14interface type interface-path-idExample:

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/3/0/0

Enters interface configuration mode and associates one or more interfaces to a different authentication type.

Step 15 authentication [message-digest | null] Example:

RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication null

Specifies no authentication on TenGigE 0/3/0/0, overriding the plain text authentication specified for area 1.

- By default, all of the interfaces configured in the same area inherit the same authentication parameter values of the area.
- Step 16 Use the commit or end command.

commit—Saves the configuration changes and remains within the configuration session.

- end —Prompts user to take one of these actions:
 - Yes Saves configuration changes and exits the configuration session.
 - No —Exits the configuration session without committing the configuration changes.
 - Cancel Remains in the configuration session, without committing the configuration changes.

Control Frequency That 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

- 1. configure
- **2.** Do one of the following:
 - router ospf process-name
 - router ospfv3 process-name
- **3.** router-id { router-id }
- **4.** Perform Step 5 or Step 6 or both to control the frequency that the same LSA is originated or accepted.

- 5. timers lsa refresh seconds
- 6. timers lsa min-arrival seconds
- 7. timers lsa group-pacing seconds
- 8. Use the commit or end command.

DETAILED STEPS

```
Step 1
           configure
           Example:
           RP/0/RP0/CPU0:router# configure
           Enters mode.
Step 2
           Do one of the following:
              • router ospf process-name
              • router ospfv3 process-name
           Example:
           RP/0/RP0/CPU0:router: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.
           Note
                       The process-name argument is any alphanumeric string no longer than 40 characters.
Step 3
           router-id { router-id }
           Example:
           RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3
           Configures a router ID for the OSPF process.
           Note
                       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
           timers lsa refresh seconds
           Example:
           RP/0/RP0/CPU0:router(config-ospf)# timers lsa refresh 1800
           Sets how often self-originated LSAs should be refreshed, in seconds.
```

• The default is 1800 seconds for both OSPF and OSPFv3.

Step 6 timers lsa min-arrival seconds

Example:

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.

• The default is 1 second.

Step 7 timers lsa group-pacing seconds

Example:

RP/0/

/CPU0:router(config-ospf)# timers lsa group-pacing 1000

Changes the interval at which OSPF link-state LSAs are collected into a group for flooding.

• The default is 240 seconds.

Step 8 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No -Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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 1: Virtual Link to Area 0

This figure illustrates a virtual link from Area 3 to Area 0.



Create Virtual Link

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.

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.

Before you begin

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. . 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.

SUMMARY STEPS

- **1.** Do one of the following:
 - **show ospf** [process-name]

• show ospfv3 [process-name]

- **2**. configure
- **3.** Do one of the following:
 - router ospf process-name
 - 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 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.
- **10.** Use the **commit** or **end** command.
- **11.** Do one of the following:
 - show ospf [process-name] [area-id] virtual-links
 - show ospfv3 [process-name] virtual-links

DETAILED STEPS

Step 1 Do one of the following:

• show ospf [process-name]

• show ospfv3 [process-name]

Example:

RP/0//CPU0:router# show ospf

or

RP/0//CPU0:router# show ospfv3

(Optional) Displays general information about OSPF routing processes.

• The output displays the router ID of the local router. You need this router ID to configure the other end of the link.

Step 2 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 3 Do one of the following:

- router ospf process-name
- router ospfv3 process-name

Example:

	RP/0//CPU	0:router(config)# router ospf 1		
	or			
	RP/0//CPU	0:router(config)# router ospfv3 1		
	Enables O	SPF routing for the specified routing process and places the router in router configuration mode.		
	or			
	Enables O	SPFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.		
	Note	The process-name argument is any alphanumeric string no longer than 40 characters.		
Step 4	router-id	{ router-id }		
	Example:			
	RP/0//CPU	0:router(config-ospf)# router-id 192.168.4.3		
	Configures	s a router ID for the OSPF process.		
	Note	We recommend using a stable IPv4 address as the router ID.		
Step 5	area arec	u-id		
	Example:			
	RP/0//CPU	0:router(config-ospf)# area 1		
	Enters area	a configuration mode and configures a nonbackbone area for the OSPF process.		
		<i>area-id</i> argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address on.		
Step 6	virtual-lin	k router-id		
	Example:			
	RRP/0//CP	U0:router(config-ospf-ar)# virtual-link 10.3.4.5		
	Defines an	OSPF virtual link.		
	• See .			
Step 7	authentica	ation message-digest		
	Example:			
	RP/0//CPU	0:router(config-ospf-ar-vl)#authentication message-digest		
	Selects MI	D5 authentication for this virtual link.		
Step 8	message-d	ligest-key key-id md5 { key clear key encrypted key }		
	Example:			
	RP/0//CPU	0:router(config-ospf-ar-vl)#message-digest-key 4 md5 vourkey		

Defines an OSPF virtual link.

- See to understand a virtual link.
- The *key-id* argument is a number in the range from 1 to 255. The *key* 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.
- The authentication-key key command is not supported for OSPFv3.
- Once the key is encrypted it must remain encrypted.
- **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** Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.
- **Step 11** Do one of the following:
 - show ospf [process-name] [area-id] virtual-links
 - show ospfv3 [process-name] virtual-links

Example:

RP/0//CPU0:router# show ospf 1 2 virtual-links

or

RP/0//CPU0:router# show ospfv3 1 virtual-links

(Optional) Displays the parameters and the current state of OSPF virtual links.

Creating virtual link- example

ABR 1 Configuration

ABR 2 Configuration

In the following example, the **show ospfv3 virtual links** 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.

```
show ospfv3 virtual-links
```

```
Virtual Links for OSPFv3 1
Virtual Link OSPF VLO 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 TenGigE 0/1/0/1 Cost of using 2
 Transmit Delay is 5 sec,
 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 VLO to router 10.0.0.3 is up
   Adjacency State FULL (Hello suppressed)
State is up and Adjacency State is FULL
```

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:

```
router ospfv3 1
router-id 10.0.0.217
area 0
interface TenGigE 0/2/0/1
area 1
virtual-link 10.0.0.212
interface TenGigE 0/2/0/0
router ospfv3 1
router-id 10.0.0.212
area 0
interface TenGigE 0/3/0/1
area 1
virtual-link 10.0.0.217
interface TenGigE 0/2/0/0
```

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 TenGigE 0/2/0/1
interface TenGigE 0/3/0/0
area 1
interface TenGigE 0/2/0/0
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 TenGigE 0/9/0/1
virtual-link 10.10.10.10
area 3
interface Loopback 0
interface TenGigE 0/9/0/0
```

Summarize Subnetwork LSAs on 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.** Do one of the following:
 - router ospf process-name
 - router ospfv3 process-name
- **3**. **router-id** { *router-id* }
- 4. area area-id
- **5.** Do one of the following:
 - range *ip-address mask* [advertise | not-advertise]
 - range *ipv6-prefix | prefix-length* [advertise | not-advertise]
- 6. interface type interface-path-id
- 7. Use the commit or end command.

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

- **Step 2** Do one of the following:
 - router ospf process-name

```
• router ospfv3 process-name
           Example:
           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.
           Note
                       The process-name argument is any alphanumeric string no longer than 40 characters.
Step 3
           router-id { router-id }
           Example:
           RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3
           Configures a router ID for the OSPF process.
                       We recommend using a stable IPv4 address as the router ID.
           Note
```

Step 4 area area-id

Example:

RP/0/RP0/CPU0:router(config-ospf)# area

Enters area configuration mode and configures a nonbackbone area for the OSPF process.

• The *area-id* 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 Do one of the following:

- range *ip-address mask* [advertise | not-advertise]
- range *ipv6-prefix | prefix-length* [advertise | not-advertise]

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# range 192.168.0.0 255.255.0.0 advertise

or

RP/0/RP0/CPU0:router(config-ospf-ar)# range 4004:f000::/32 advertise

Consolidates and summarizes OSPF routes at an area boundary.

- The **advertise** keyword causes the software to advertise the address range of subnetworks in a Type 3 summary LSA.
- The **not-advertise** keyword causes the software to suppress the Type 3 summary LSA, and the subnetworks in the range remain hidden from other areas.

- In the first example, all subnetworks for network 192.168.0.0 are summarized and advertised by the ABR into areas outside the backbone.
- In the second example, two or more IPv4 interfaces are covered by a 192.x.x network.

Step 6 interface type interface-path-id

```
Example:
```

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/0/0/0

Enters interface configuration mode and associates one or more interfaces to the area.

Step 7 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

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 TenGigE 0/0/0/0
area 1
range 2300::/16
interface TenGigE 0/0/0/0
```

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.

Redistribute Routes into OSPF

This task redistributes routes from an IGP (could be a different OSPF process) into OSPF.

SUMMARY STEPS

- 1. configure
- **2.** Do one of the following:

- router ospf process-name
- router ospfv3 process-name
- **3.** router-id { router-id }
- **4.** redistribute *protocol* [*process-id*] { level-1 | level-2 } [metric *metric-value*] [metric-type *type-value*] [match { external [1 | 2]} [tag *tag-value*] [route-policy *policy-name*]
- **5.** Do one of the following:
 - summary-prefix address mask [not-advertise][tag tag]
 - summary-prefix ipv6-prefix | prefix-length [not-advertise] [tag tag]
- 6. Use the commit or end command.

DETAILED STEPS

configure		
Example:		
RP/0/RP0	/CPU0:router# configure	
Enters mo	ide.	
Do one of	the following:	
	er ospf process-name er ospfv3 process-name	
Example:		
RP/0/RP0	/CPU0:router(config)# router ospf 1	
or		
RP/0/RP0/CPU0:router(config) # router ospfv3 1		
Enables C	SPF 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.		
Note	The process-name argument is any alphanumeric string no longer than 40 characters.	
router-id	{ router-id }	
Example:		
RRP/0/RP	0/CPU0:router(config-ospf)# router-id 192.168.4.3	
Configure	es a router ID for the OSPF process.	

RP/0/RP0/CPU0:router(config-ospf)# redistribute bgp 100

or

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

Redistributes OSPF routes from one routing domain to another routing domain.

or

Redistributes OSPFv3 routes from one routing domain to another routing domain.

- This command causes the router to become an ASBR by definition.
- OSPF tags all routes learned through redistribution as external.
- The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.
- 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.
- The OSPF example redistributes BGP autonomous system 1, Level 1 routes into OSPF as Type 2 external routes.
- 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.

Note RPL is not supported for OSPFv3.

Step 5 Do one of the following:

summary-prefix address mask [not-advertise][tag tag]
summary-prefix ipv6-prefix | prefix-length [not-advertise][tag tag]

Example:

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

(Optional) Creates aggregate addresses for OSPF.

or

(Optional) Creates aggregate addresses for OSPFv3.

- This command provides external route summarization of the non-OSPF routes.
- 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.
- This command is optional. If you do not specify it, each route is included in the link-state database and advertised in LSAs.
- 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.
- 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.

Step 6 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- **Cancel**—Remains in the configuration session, without committing the configuration changes.

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 TenGigE 0/2/0/0
```

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 2: SPF Calculation Intervals Set by the timers spf Command

This figure 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.




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 3: Timer Intervals Reset After Topology Change Event, on page 60 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 3: Timer Intervals Reset After Topology Change Event



Configure 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.

SUMMARY STEPS

- 1. configure
- **2.** Do one of the following:
 - router ospf process-name
 - 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 interface-path-id
- 7. Use the commit or end command.
- **8.** Do one of the following:

• show ospf [process-name]

• show ospfv3 [process-name]

DETAILED STEPS

I

Step 1	configure			
	Example:			
	RP/0/RP0/CPU0:router# configure			
	Enters mod	e.		
Step 2	Do one of t	he following:		
	 router ospf process-name router ospfv3 process-name 			
	Example:			
	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 OS	PFv3 routing for the specified routing process and places the router in router ospfv3 configuration mode.		
	Note	The process-name argument is any alphanumeric string no longer than 40 characters.		
Step 3	Step 3 router-id { router-id }			
	Example:			
RP/0/RP0/CPU0:router(config-ospf)# router-id 192.168.4.3		CPU0:router(config-ospf)# router-id 192.168.4.3		
	Configures	a router ID for the OSPF process.		
	Note	We recommend using a stable IPv4 address as the router ID.		
Step 4	timers thro	ottle spf spf-start spf-hold spf-max-wait		
	Example:			
	RP/0/RP0/0	CPU0:router(config-ospf)# timers throttle spf 10 4800 90000		
	Sets SPF th	rottling timers.		
Step 5	area area-	-id		
	Example:			
	RP/0/RP0/0	CPU0:router(config-ospf)# area 0		

Enters area configuration mode and configures a backbone area.

• The *area-id* 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 6 interface type interface-path-id

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/0/0/0

Enters interface configuration mode and associates one or more interfaces to the area.

Step 7 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No --Exits the configuration session without committing the configuration changes.
- **Cancel**—Remains in the configuration session, without committing the configuration changes.

Step 8 Do one of the following:

- show ospf [process-name]
- show ospfv3 [process-name]

Example:

```
RP/0/RP0/CPU0:router# show ospf 1
```

or

RP/0/RP0/CPU0:router# RP/0/RP0/CPU0:router# show ospfv3 2

(Optional) Displays SPF throttling timers.

Graceful Restart for OSPFv3

The OSPFv3 Graceful Shutdown feature preserves the data plane capability in these circumstances:

- Planned OSPFv3 process restart, such as a restart resulting from a software upgrade or downgrade
- · Unplanned OSPFv3 process restart, such as a restart resulting from a process crash

In addition, OSPFv3 will unilaterally shutdown and enter the exited state when a critical memory event, indicating the processor is critically low on available memory, is received from the sysmon watch dog process.

This feature supports nonstop data forwarding on established routes while the OSPFv3 routing protocol restarts. Therefore, this feature enhances high availability of IPv6 forwarding.

Configure 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 disable
- 7. Use the commit or end command.
- 8. show ospfv3 [process-name [area-id]] database grace

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router ospfv3 process-name
	Example:
	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	graceful-restart
	Example:
	RP/0/RP0/CPU0:router(config-ospfv3)#graceful-restart
	Enables graceful restart on the current router.
Step 4	graceful-restart lifetime
	Example:
	RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart lifetime 120
	Specifies a maximum duration for a graceful restart.
	• The default lifetime is 95 seconds.
	• The range is 90 to 3600 seconds.
Step 5	graceful-restart interval seconds

Example:

RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart interval 120

Specifies the interval (minimal time) between graceful restarts on the current router.

- The default value for the interval is 90 seconds.
- The range is 90 to 3600 seconds.

Step 6 graceful-restart helper disable Example:

RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart helper disable
Disables the helper capability.

Step 7 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

 Step 8
 show ospfv3 [process-name [area-id]] database grace

Example:

RP/0/RP0/CPU0:router# show ospfv3 1 database grace

Displays the state of the graceful restart link.

Display 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 msecs Minimum hold time between two consecutive SPFs 10000 msecs Maximum wait time between two consecutive SPFs 10000 msecs

Initial LSA throttle delay 0 msecs Minimum hold time for LSA throttle 5000 msecs Maximum wait time for LSA throttle 5000 msecs Minimum LSA arrival 1000 msecs LSA group pacing timer 240 secs Interface flood pacing timer 33 msecs Retransmission pacing timer 66 msecs Maximum number of configured interfaces 255 Number of external LSA 0. Checksum Sum 0000000 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

OSPFv2OSPF SPF Prefix Prioritization

The OSPFv2 OSPF SPF Prefix Prioritization feature enables an administrator to converge, in a faster mode, important prefixes during route installation.

When a large number of prefixes must be installed in the Routing Information Base (RIB) and the Forwarding Information Base (FIB), the update duration between the first and last prefix, during SPF, can be significant.

In networks where time-sensitive traffic (for example, VoIP) may transit to the same router along with other traffic flows, it is important to prioritize RIB and FIB updates during SPF for these time-sensitive prefixes.

The OSPFv2OSPF SPF Prefix Prioritization feature provides the administrator with the ability to prioritize important prefixes to be installed, into the RIB during SPF calculations. Important prefixes converge faster among prefixes of the same route type per area. Before RIB and FIB installation, routes and prefixes are assigned to various priority batch queues in the OSPF local RIB, based on specified route policy. The RIB priority batch queues are classified as "critical," "high," "medium," and "low," in the order of decreasing priority.

When enabled, prefix alters the sequence of updating the RIB with this prefix priority:

Critical > High > Medium > Low

As soon as prefix priority is configured, /32 prefixes are no longer preferred by default; they are placed in the low-priority queue, if they are not matched with higher-priority policies. Route policies must be devised to retain /32s in the higher-priority queues (high-priority or medium-priority queues).

Priority is specified using route policy, which can be matched based on IP addresses or route tags. During SPF, a prefix is checked against the specified route policy and is assigned to the appropriate RIB batch priority queue.

These are examples of this scenario:

- If only high-priority route policy is specified, and no route policy is configured for a medium priority:
 - Permitted prefixes are assigned to a high-priority queue.
 - Unmatched prefixes, including /32s, are placed in a low-priority queue.

- If both high-priority and medium-priority route policies are specified, and no maps are specified for critical priority:
 - Permitted prefixes matching high-priority route policy are assigned to a high-priority queue.
 - Permitted prefixes matching medium-priority route policy are placed in a medium-priority queue.
 - Unmatched prefixes, including /32s, are moved to a low-priority queue.
- If both critical-priority and high-priority route policies are specified, and no maps are specified for medium priority:
 - · Permitted prefixes matching critical-priority route policy are assigned to a critical-priority queue.
 - Permitted prefixes matching high-priority route policy are assigned to a high-priority queue.
 - Unmatched prefixes, including /32s, are placed in a low-priority queue.
- If only medium-priority route policy is specified and no maps are specified for high priority or critical priority:
 - Permitted prefixes matching medium-priority route policy are assigned to a medium-priority queue.
 - Unmatched prefixes, including /32s, are placed in a low-priority queue.

Use the **[no] spf prefix-priority route-policy** *rpl* command to prioritize OSPFv2OSPF prefix installation into the global RIB during SPF.

SPF prefix prioritization is disabled by default. In disabled mode, /32 prefixes are installed into the global RIB, before other prefixes. If SPF prioritization is enabled, routes are matched against the route-policy criteria and are assigned to the appropriate priority queue based on the SPF priority set. Unmatched prefixes, including /32s, are placed in the low-priority queue.

If all /32s are desired in the high-priority queue or medium-priority queue, configure this single route map:

```
prefix-set ospf-medium-prefixes
   0.0.0.0/0 ge 32
   end-set
```

Configure OSPFv2 OSPF SPF Prefix Prioritization

Perform this task to configure OSPFv2 OSPF SPF (shortest path first) prefix prioritization.

SUMMARY STEPS

- 1. configure
- 2. prefix-set prefix-set name
- **3.** route-policy *route-policy name* if destination in *prefix-set name* then set spf-priority {critical | high | medium} endif
- **4.** Use one of these commands:
 - router ospf ospf-name
 - router ospfv3 ospfv3-name

- **5.** router ospf ospf name
- 6. spf prefix-priority route-policy route-policy name
- 7. Use the commit or end command.
- 8. show rpl route-policy route-policy name detail

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 2 prefix-set prefix-set name

Example:

RP/0/RP0/CPU0:router(config)#prefix-set ospf-critical-prefixes RP/0/RP0/CPU0:router(config-pfx)#66.0.0.0/16 RP/0/RP0/CPU0:router(config-pfx)#end-set

Configures the prefix set.

Step 3 route-policy *route-policy name* if destination in *prefix-set name* then set spf-priority {critical | high | medium} endif

Example:

```
RP/0/RP0/CPU0:router#route-policy ospf-spf-priority
RP/0/RP0/CPU0:router(config-rpl)#if destination in ospf-critical-prefixes then
set spf-priority critical
endif
RP/0/RP0/CPU0:router(config-rpl)#end-policy
```

Configures route policy and sets OSPF SPF priority.

Step 4 Use one of these commands:

- router ospf ospf-name
- router ospfv3 ospfv3-name

Example:

```
RP/0/RP0/CPU0:router# router ospf 1
```

Or

RP/0/RP0/CPU0:router# router ospfv3 1

Enters Router OSPF configuration mode.

Step 5 router ospf ospf name

Example:

RP/0/RP0/CPU0:router# router ospf 1

Enters Router OSPF configuration mode.

 Step 6
 spf prefix-priority route-policy route-policy name

 Example:

RP/0/RP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority

Or

RP/0/RP0/CPU0:router(config-ospfv3)#spf prefix-priority route-policy ospf3-spf-priority Configures SPF prefix-priority for the defined route policy.

Note Configure the **spf prefix-priority** command under router OSPF.

Step 7 Use the **commit** or **end** command.

commit —Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

Step 8 show rpl route-policy route-policy name detail

Example:

```
RP/0/RP0/CPU0:router#show rpl route-policy ospf-spf-priority detail
  prefix-set ospf-critical-prefixes
     66.0.0.0/16
  end-set
  !
  route-policy ospf-spf-priority
    if destination in ospf-critical-prefixes then
        set spf-priority critical
    endif
  end-policy
  !
```

Displays the set SPF prefix priority.

OSPFv2

OSPFv3

This example shows how to configure /32 prefixes as medium-priority, in general, in addition to placing some /32 and /24 prefixes in critical-priority and high-priority queues:

```
prefix-set ospf-critical-prefixes
```

```
192.41.5.41/32,
  11.1.3.0/24,
 192.168.0.44/32
  end-set
  1
 prefix-set ospf-high-prefixes
  44.4.10.0/24,
  192.41.4.41/32,
 41.4.41.41/32
 end-set
  1
 prefix-set ospf-medium-prefixes
 0.0.0.0/0 ge 32
 end-set
  !
  route-policy ospf-priority
   if destination in ospf-high-prefixes then
     set spf-priority high
   else
      if destination in ospf-critical-prefixes then
        set spf-priority critical
      else
       if destination in ospf-medium-prefixes then
         set spf-priority medium
        endif
      endif
   endif
  end-policy
router ospf 1
   spf prefix-priority route-policy ospf-priority
   area O
   interface TenGigE 0/3/0/0
    !
   1
   area 3
   interface TenGigE 0/2/0/0
   1
   !
   area 8
   interface TenGigE 0/2/0/0
router ospfv3 1
   spf prefix-priority route-policy ospf-priority
   area O
   interface TenGigE 0/3/0/0
   1
   !
   area 3
   interface TenGigE 0/2/0/0
    !
   1
   area 8
   interface TenGigE 0/2/0/0
```

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. You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.
- 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.

Configure Multi-area Adjacency

This task explains how to create multiple areas on an OSPF primary interface.

Before you begin



Note You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

SUMMARY STEPS

- **1**. configure
- 2. router ospf process-name
- 3. area area-id
- **4. interface** *type interface-path-id*
- 5. area area-id
- 6. multi-area-interface type interface-path-id
- 7. Use the commit or end command.

DETAILED STEPS

Step 1	configure		
	Example:		
	RP/0/RP0/CPU0:router# configure		
	Enters mode.		
Step 2	router ospf process-name		
	Example:		
	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.		
	Note The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.		
Step 3	area area-id		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf)# area 0		
	Enters area configuration mode and configures a backbone area.		
	• 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 4	interface type interface-path-id		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf-ar)# interface Serial 0/1/0/3		
	Enters interface configuration mode and associates one or more interfaces to the area.		
Step 5	area area-id		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf)# area 1		
	Enters area configuration mode and configures an area used for multiple area adjacency.		
	• 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 6	multi-area-interface type interface-path-id		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf)# multi-area-interface Serial 0/1/0/3		
	Enables multiple adjacencies for different OSPF areas and enters multi-area interface configuration mode		

Step 7 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel —Remains in the configuration session, without committing the configuration changes.

Example

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:

<pre>config-ospf-ar) # multi-area-interface TenGigE 0/0/0/0</pre>
config-ospf-ar-mif)# ?
Enable authentication
Authentication password (key)
Interface cost
Cost when cumulative bandwidth goes below the theshold
Filter OSPF LSA during synchronization and flooding
Interval after which a neighbor is declared dead
Filter networks in routing updates
Time between HELLO packets
Message digest authentication password (key)
Enable/Disable ignoring of MTU in DBD packets
Customize size of OSPF packets upto MTU
Time between retransmitting lost link state advertisements
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 **igp auto-config disable** command. This allows LDP to receive all OSPF interfaces minus the ones explicitly disabled.

Configure 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

- 1. configure
- 2. router ospf process-name
- **3**. mpls ldp auto-config
- 4. Use the commit or end command.

DETAILED STEPS

Step 1	configure		
	Example:		
	RP/0/RP0/CPU0:router# configure		
	Enters mode.		
Step 2	router ospf process-name		
	Example:		
	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.		
	Note The <i>process-name</i> argument is any alphanumeric string no longer than 40 characters.		
Step 3	mpls ldp auto-config		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf)# mpls ldp auto-config		
	Enables LDP IGP interface auto-configuration for an OSPF instance.		
	• Optionally, this command can be configured for an area of an OSPF instance.		
Step 4	Use the commit or end command.		
	commit —Saves the configuration changes and remains within the configuration session.		
	end —Prompts user to take one of these actions:		
	• Yes — Saves configuration changes and exits the configuration session.		
	• No —Exits the configuration session without committing the configuration changes.		
	• Cancel — Remains in the configuration session, without committing the configuration changes.		

Configure LDP IGP Synchronization: OSPF

Perform this task to configure LDP IGP Synchronization under OSPF.



By default, there is no synchronization between LDP and IGPs.

SUMMARY STEPS

- 1. configure
- 2. router ospf process-name
- **3.** (Optional) **vrf** *vrf*-name
- **4.** Use one of the following commands:
 - mpls ldp sync
 - area area-id mpls ldp sync
 - area area-id interface name mpls ldp sync
- **5.** (Optional) Use one of the following commands:
 - mpls ldp sync
 - area area-id mpls ldp sync
 - area area-id interface name mpls ldp sync
- 6. Use the commit or end command.
- 7. (Optional) show mpls ldp vrf vrf-name ipv4 igp sync
- 8. (Optional) show mpls ldp vrf all ipv4 igp sync
- **9.** (Optional) show mpls ldp { ipv4 | ipv6 } igp sync

DETAILED STEPS

Step 1	configure Example:		
	RP/0/RP0/CPU0:router# configure		
	Enters mode.		
Step 2	router ospf process-name		
	Example:		
	RP/0/RP0/CPU0:router(config)# router ospf 100		
	Identifies the OSPF routing process and enters OSPF configuration mode.		
Step 3	(Optional) vrf vrf-name		
	Example:		
	RP/0/RP0/CPU0:router(config-ospf) # vrf red		

I

	Specifies the non-default VRF.
Step 4	Use one of the following commands:
	• mpls ldp sync
	• area area-id mpls ldp sync
	area area-id interface name mpls ldp sync Example:
	RP/0/RP0/CPU0:router(config-ospf) # mpls ldp sync Enables LDP IGP synchronization on an interface.
Ston E	
Step 5	(Optional) Use one of the following commands:
	 mpls ldp sync area area-id mpls ldp sync
	• area area-id interface name mpls ldp sync
	Example:
	RP/0/RP0/CPU0:router(config-ospf-vrf) # mpls ldp sync
	<pre>RP/0/RP0/CPU0:router(config-ospf-vrf)# area 1 mpls ldp sync</pre>
	Enables LDP IGP synchronization on an interface for the specified VRF.
Step 6	Use the commit or end command.
	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.
	• No —Exits the configuration session without committing the configuration changes.
	• Cancel — Remains in the configuration session, without committing the configuration changes.
Step 7	(Optional) show mpls ldp vrf vrf-name ipv4 igp sync
	Example:
	RP/0/RP0/CPU0:router# show mpls ldp vrf red ipv4 igp sync
	Displays the LDP IGP synchronization information for the specified VRF for address family IPv4.
Step 8	(Optional) show mpls ldp vrf all ipv4 igp sync
	Example:
	RP/0/RP0/CPU0:router# show mpls ldp vrf all ipv4 igp sync
	Displays the LDP IGP synchronization information for all VRFs for address family IPv4.
Step 9	(Optional) show mpls ldp { ipv4 ipv6 } igp sync
-	

Example:

RP/0/RP0/CPU0:router# show mpls ldp ipv4 igp sync

RP/0/RP0/CPU0:router# show mpls ldp ipv6 igp sync

Displays the LDP IGP synchronization information for IPv4 or IPv6 address families.

Example

The example shows how to configure LDP IGP synchronization for OSPF.

```
router ospf 100
mpls ldp sync
!
mpls ldp
igp sync delay 30
!
```

IGP link state

IGP Link-State Database Distribution

A given BGP node may have connections to multiple, independent routing domains. IGP link-state database distribution into BGP-LS is supported for both OSPF and IS-IS protocols in order to distribute this information on to controllers or applications that desire to build paths spanning or including these multiple domains.

To distribute OSPFv2 link-state data using BGP-LS, use the **distribute link-state** command in router configuration mode.

```
Router# configure
Router(config)# router ospf 100
Router(config-ospf)# distribute link-state instance-id 32
```

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.

Configure Authentication Message Digest Management for OSPF

This task explains how to manage authentication of a keychain on the OSPF interface.

Before you begin

A valid keychain must be configured before this task can be attempted.

SUMMARY STEPS

- 1. configure
- **2.** router ospf process-name
- **3.** router-id { router-id }
- 4. area area-id
- **5. interface** *type interface-path-id*
- 6. authentication [message-digest keychain | null]
- 7. Use the commit or end command.

DETAILED STEPS

Step 1	configure		
	Example:		
	RP/0/RP0/C	PU0:router# configure	
	Enters mode	2.	
Step 2	process-name		
	Example:		
	RP/0/RP0/C	PU0:router(config)# router ospf 1	
	Enables OS	PF routing for the specified routing process and places the router in router configuration mode.	
	Note	The process-name argument is any alphanumeric string no longer than 40 characters.	
Step 3	3 router-id { router-id }		
	Example:		
	RP/0/RP0/C	PU0:router(config-ospf)# router id 192.168.4.3	
	Configures a router ID for the OSPF process.		
	Note	We recommend using a stable IPv4 address as the router ID.	
Step 4	area area-	id	

Example:

RP/0/RP0/CPU0:router(config-ospf) # area 1

Enters area configuration mode.

The *area-id* 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 interface type interface-path-id

Example:

RP/0/RP0/CPU0:router(config-ospf-ar)# interface TenGigE 0/0/0/0

Enters interface configuration mode and associates one or more interfaces to the area.

Step 6 authentication [message-digest keychain | null]

Configures an MD5 keychain.

Example:

The following example shows the configuration for message-digest authentication.

RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication message-digest keychain ospf_int1

Note In the above example, the *ospf_intl* keychain must be configured before you attempt this step.

Step 7 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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 MD5
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 TenGigE 0/0/0/0 interface:

show ospf 1 interface TenGigE 0/0/0/0

```
TenGigE 0/0/0/0 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:

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

OSPF Link-State Database Overload Protection

The OSPF Link-State Database Overload Protection feature allows you to protect the OSPF routing process from the large number of received link-state advertisements (LSAs) that can result from a misconfiguration on another router in the OSPF domain (for example, the redistribution of a large number of IP prefixes to OSPF). In this feature, the router monitors the count of all nonself-generated LSAs it receives. When the total number of LSAs reaches the value set as a threshold in the configuration, the router logs a warning message. When the total LSAs exceed the max number in configuration, the router stops accepting new LSAs. Therefore, the OSPF Link-State Database Overload Protection feature allows you to limit the number of nonself-generated LSAs for a given OSPF process. Excessive LSAs generated by other routers in the OSPF domain can substantially drain the CPU and memory resources of the router.

If the count of received LSAs is higher than the configured max number after one minute, the OSPF process disables all adjacencies in the given context and clears the OSPF database. This state is called the ignore state. In this state, the router ignores all OSPF packets received on all interfaces belonging to the OSPF instance and stops OSPF packet generation on its interfaces. The OSPF process remains in ignore state for the duration of the configured ignore time. When the ignore time expires, the OSPF process returns to a normal operation state and starts building adjacencies on all its interfaces.

To prevent the OSPF instance from endlessly oscillating between the normal state and ignore state, as a result of the LSA counts immediately exceeding the max number again after it returns from the ignore state, the OSPF instance keeps a count of how many times it has been in ignore state. This counter is called the ignore count. If the ignore count exceeds its configured value, the OSPF instance remains in ignore state permanently.

To return the OSPF instance to its normal state, use the **clear IP OSPF process** command. The ignore count is reset to zero if the LSA count doesn't exceed the max number again during the time configured by the reset-time keyword.



Note

The are no default values for the **max-lsa** command. The router considers the max-lsa limits only if it's available in the configuration. After configuring the OSPF Link-State Database Overload Protection feature, the default values for other OSPF parameters are as follows:

- Default threshold percentage to log warning: 75%.
- Default ignore count value: 5
- Default ignore time in minutes: 5 minutes
- Default time to reset ignore count: 10 minutes

In the event of LSA count reaches or exceeds the configured threshold value, router displays the following logs:

```
Router:Oct 25 11:42:49.756 IST: ospf[1046]: %ROUTING-OSPF-4-MAX_LSA_THR : Threshold for
maximum number of non self-generated LSAs has been reached "default" - 9000 LSAs
Router:Oct 25 11:42:49.756 IST: ospf[1046]: %ROUTING-OSPF-4-MAX_LSA : Maximum number of non
self-generated LSAs has been exceeded "default" - 12001 LSAs
```

Router:Oct 25 11:42:49.756 IST: ospf[1046]: %ROUTING-OSPF-4-MAX LSA THR : Threshold for

```
maximum number of non self-generated LSAs has been reached "V1" - 750 LSAs
Router:Oct 25 11:42:49.756 IST: ospf[1046]: %ROUTING-OSPF-4-MAX_LSA : Maximum number of non
self-generated LSAs has been exceeded "V1" - 1001 LSA
```

Configuration Example

The following example shows how to configure the OSPF instance to accept 12000 nonself-generated LSAs in the global routing table and 1000 nonself-generated LSAs in VRF V1:

```
Router# configure
Router(config)# router ospf 0
Router(config-ospf)# max-lsa 12000
Router(config-ospf)# vrf V1
Router(config-ospf)# max-lsa 1000
```

Running Configuration

The following example shows how to display the current status of the OSPF instance:

```
Router# show ospf 0
  Routing Process "ospf 0" with ID 10.0.0.2
  NSR (Non-stop routing) is Disabled
 Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
 Maximum number of non self-generated LSA allowed 12000
   Current number of non self-generated LSA 1
    Threshold for warning message 75%
    Ignore-time 5 minutes, reset-time 10 minutes
   Ignore-count allowed 5, current ignore-count 0
Router# show ospf vrf V1
VRF V1 active in Routing Process "ospf 0" with ID 10.0.0.2 \,
Role: Primary Active
NSR (Non-stop routing) is Enabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
Maximum number of non self-generated LSA allowed 1000
    Current number of non self-generated LSA 1
   Threshold for warning message 75%
   Ignore-time 5 minutes, reset-time 10 minutes
   Ignore-count allowed 5, current ignore-count 0
Router is not originating router-LSAs with maximum metric
```

References for OSPF

To implement OSPF you need to understand the following concepts:

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.

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.

Comparison of Cisco IOS XR Software 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 Software 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

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 the software, 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.

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 4: OSPF Routing Components

This figure illustrates the routing components in an OSPF network topology.



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: OSPF Routing Components* 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 the *Figure 1: OSPF Routing Components*, 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: OSPF Routing Components*, Area 0 is an OSPF backbone area. Any OSPF backbone area has a reserved area ID of 0.0.0.

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: OSPF Routing Components* section, 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

An interior router (such as R1 in *Figure 1: OSPF Routing Components*) is 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 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. OSPF obtains a router ID from the following sources, in order of decreasing preference:

• By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.

- 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 ITAL selected router-id.
- 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
- Broadcast networks

You can configure your 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).

OSPV 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. See the *OSPFAuthentication Message Digest Management* section.

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.

OSPF FIB Download Notification

OSPF FIB Download Notification feature minimizes the ingress traffic drop for a prolonged period of time after the line card reloads.

Open Shortest Path First (OSPF) registers with Routing Information Base (RIB) through ITAL which keeps the interface down until all the routes are downloaded to Forwarding Information Base (FIB). OSPF gets the Interface Up notification when all the routes on the reloaded line card are downloaded through RIB/FIB.

RIB provides notification to registered clients when a:

- Node is lost.
- Node is created.
- Node's FIB upload is completed.

Designated Router (DR) for OSPF

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.

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.

Passive Interface

Setting an interface as passive disables the sending of routing updates for the neighbors, hence adjacencies will not be formed in OSPF. However, the particular subnet will continue to be advertised to OSPF neighbors. Use the **passive** command in appropriate mode to suppress the sending of OSPF protocol operation on an interface.

It is recommended to use passive configuration on interfaces that are connecting LAN segments with hosts to the rest of the network, but are not meant to be transit links between routers.

Modes of Graceful Restart Operation

The operational modes that a router can be in for this feature are restart mode and helper moderestart mode, helper mode, and protocol shutdown 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.

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.

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 OSFPv3 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.

Protocol Shutdown Mode

In this mode the OSPFv3 operation is completely disabled. This is accomplished by flushing self-originated link state advertisements (LSAs), immediately bringing down local OSPFv3-supported interfaces, and clearing the Link State Database (LSDB). The non-local LSDB entries are removed by OSPFv3, These are not flooded (MaxAged).

The protocol shutdown mode can be invoked either manually through the **protocol shutdown** command that disables the protocol instance or when the OSPFv3 process runs out of memory. These events occur when protocol shut down is performed:

- The local Router LSA and all local Link LSAs are flushed. All other LSAs are eventually aged out by other OSPFv3 routers in the domain.
- OSPFv3 neighbors not yet in Full state with the local router are brought down with the Kill_Nbr event.
- After a three second delay, empty Hello packets are immediately sent to each neighbor that has an active adjacency.
 - An empty Hello packet is sent periodically until the dead interval has elapsed.
 - When the dead_interval elapses, Hello packets are no longer sent.

After a Dead Hello interval delay (4 X Hello Interval), the following events are then performed:

- The LSA database from that OSPFv3 instance is cleared.
- All routes from RIB that were installed by OSPFv3 are purged.

The router will not respond to any OSPF control packets it receives from neighbors while in protocol shutdown state.

Protocol Restoration

The method of restoring the protocol is dependent on the trigger that originally invoked the shut down. If the OSPFv3 was shut down using the **protocol shutdown** command, then use the **no protocol shutdown** command to restore OSPFv3 back to normal operation. If the OSPFv3 was shutdown due to a Critical Memory message from the sysmon, then a Normal Memory message from sysmon, which indicates that sufficient memory has been restored to the processor, restores the OSPFv3 protocol to resume normal operation. When OSPFv3 is shutdown due to the Critical Memory trigger, it must be manually restarted when normal memory levels are restored on the route processor. It will not automatically restore itself.

These events occur when the OSPFv3 is restored:

- 1. All OSPFv3 interfaces are brought back up using the Hello packets and database exchange.
- 2. The local router and link LSAs are rebuilt and advertised.
- 3. The router replies normally to all OSPFv3 control messages received from neighbors.

4. Routes learned from other OSPFv3 routers are installed in RIB.

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 range for maximum paths is from 1 to 8 and the default number of maximum paths is 8.

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 reoriginates 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* module of the *MPLS Configuration guide* and the following IETF drafts:

- draft-ietf-ospf-cap-09
- draft-ietf-pce-disco-proto-ospf-00

Management Information Base (MIB) for OSPFv3

Cisco IOS XR supports full MIBs and traps for OSPFv3, as defined in RFC 5643. The RFC 5643 defines objects of the Management Information Base (MIB) for use with the Open Shortest Path First (OSPF) Routing Protocol for IPv6 (OSPF version 3).

The OSPFv3 MIB implementation is based on the IETF draft *Management Information Base for OSPFv3 (draft-ietf-ospf-ospfv3-mib-8)*. Users need to update the NMS application to pick up the new MIB when upgraded to RFC 5643.

Multiple OSPFv3 Instances

SNMPv3 supports "contexts" that can be used to implement MIB views on multiple OSPFv3 instances, in the same system.

OSPFv3 Timers Update

The Open Shortest Path First version 3 (OSPFv3) timers link-state advertisements (LSAs) and shortest path first (SPF) throttle default values are updated to:

- timers throttle lsa all-start-interval: 50 milliseconds and hold-interval: 200 milliseconds
- timers throttle spf *spf-start*: 50 milliseconds, *spf-hold*: 200 milliseconds, *spf-max-wait*: 5000 milliseconds

I



Implementing and Monitoring RIB

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.

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.

This module describes how to implement and monitor RIB on your network.

- Verify RIB Configuration Using Routing Table, on page 95
- Verify Networking and Routing Problems, on page 96
- Disable RIB Next-hop Dampening, on page 98
- Enable RCC and LCC On-demand Scan, on page 99
- Enable RCC and LCC Background Scan, on page 100
- References for RIB, on page 102

Verify RIB Configuration Using 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 | safi-all] summary [detail][standby]
- **2.** show route [vrf {vrf-name | all }][afi-all | ipv4 | ipv6][unicast | safi-all][protocol [instance]| ip-address mask][standby][detail]

DETAILED STEPS

Step 1 show route [vrf { vrf-name | all }] [afi-all | ipv4 | ipv6] [unicast | safi-all] summary [detail] [standby]

Example:

RP/0/RP0/CPU0:router# show route summary

Displays route summary information about the specified routing table.

- The default table summarized is the IPv4 unicast routing table.
- **Step 2 show route** [**vrf** {*vrf-name* | **all** }][**afi-all** | **ipv4** | **ipv6**][**unicast** | **safi-all**][*protocol* [*instance*]| *ip-address mask*][**standby**][**detail**]

Example:

RP/0/RP0/CPU0:router# show route ipv4 unicast

Displays more detailed route information about the specified routing table.

• 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.

Output of show route best-local Command: Example

The following is sample output from the **show route backup** command:

show route backup

```
Codes: C - connected, S - static, 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
S 172.73.51.0/24 is directly connected, 2d20h, HundredGigE 4/0/0/1
Backup O E2 [110/1] via 10.12.12.2, HundredGigE 3/0/0/1
```

Verify 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 | safi-all] [protocol [instance] | ip-address mask] [standby] [detail]
- **2.** show route [vrf {vrf-name | all }][afi-all | ipv4 | ipv6][unicast | safi-all] backup [ip-address][standby]
- **3.** show route [vrf {vrf-name | all }][ipv4 | ipv6][unicast | safi-all] best-local *ip-address* [standby]

- 4. show route [vrf {vrf-name | all }] [afi-all | ipv4 | ipv6] [unicast | safi-all] connected [standby 1
- 5. show route [vrf { vrf-name | all }] [afi-all | ipv4 | ipv6] [unicast | safi-all] local [interface] [standby]
- 6. show route [vrf {vrf-name | all }][ipv4 | ipv6][unicast | safi-all] longer-prefixes { *ip-address mask* | *ip-address / prefix-length* } [**standby**]
- 7. show route [vrf { vrf-name | all }] [ipv4 | ipv6] [unicast | safi-all] next-hop ip-address [standby]

DETAILED STEPS

Step 1	show route [vrf { <i>vrf-name</i> all }] [afi-all ipv4 ipv6] [unicast safi-all] [<i>protocol</i> [<i>instance</i>] <i>ip-address</i> mask] [standby] [detail]
	Example:
	RP/0/RP0/CPU0:router# show route ipv4 unicast 192.168.1.11/8
	Displays the current routes in RIB.
Step 2	<pre>show route [vrf { vrf-name all }] [afi-all ipv4 ipv6] [unicast safi-all] backup [ip-address] [standby]</pre>
	Example:
	RP/0/RP0/CPU0:router# show route ipv4 unicast backup 192.168.1.11/8
	Displays backup routes in RIB.
Step 3	<pre>show route [vrf { vrf-name all }] [ipv4 ipv6] [unicast safi-all] best-local ip-address [standby] Example:</pre>
	RP/0/RP0/CPU0:router# show route ipv4 unicast best-local 192.168.1.11/8
	Displays the best-local address to use for return packets from the given destination.
Step 4	show route [vrf {vrf-name all }][afi-all ipv4 ipv6][unicast safi-all] connected [standby]
	Example:
	RP/0/RP0/CPU0:router# show route ipv4 unicast connected
	Displays the current connected routes of the routing table.
Step 5	show route [vrf {vrf-name all }] [afi-all ipv4 ipv6] [unicast safi-all] local [interface] [standby]
	Example:
	RP/0/RP0/CPU0:router# show route ipv4 unicast local
	Displays local routes for receive entries in the routing table.
Step 6	<pre>show route [vrf { vrf-name all }] [ipv4 ipv6] [unicast safi-all] longer-prefixes { ip-address mask ip-address prefix-length } [standby]</pre>
	Example:

RP/0/RP0/CPU0:router# show route ipv4 unicast longer-prefixes 192.168.1.11/8

Displays the current routes in RIB that share a given number of bits with a given network.

 Step 7
 show route [vrf {vrf-name | all }] [ipv4 | ipv6] [unicast | safi-all] next-hop ip-address [standby]

 Example:

RP/0/RP0/CPU0:router# show route ipv4 unicast next-hop 192.168.1.34

Displays the next-hop gateway or host to a destination address.

Output of show route Command: Example

The following is sample output from the **show route** command when entered without an address:

show route

```
Codes: C - connected, S - static, 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, \star - 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
     10.2.210.0/24 is directly connected, 1d21h, Ethernet0/1/0/0
С
L
    10.2.210.221/32 is directly connected, 1d21h, Ethernet0/1/1/0
С
    172.20.16.0/24 is directly connected, 1d21h, ATM4/0.1
    172.20.16.1/32 is directly connected, 1d21h, ATM4/0.1
Τ.
С
    10.6.100.0/24 is directly connected, 1d21h, Loopback1
    10.6.200.21/32 is directly connected, 1d21h, Loopback0
L
    192.168.40.0/24 [1/0] via 172.20.16.6, 1d21h
S
```

Disable 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.** Use the **commit** or **end** command.

DETAILED STEPS

Step 1 router rib

Example:

RP/0/RP0/CPU0:router# router rib

Enters RIB configuration mode.

 Step 2
 address-family { ipv4 | ipv6 } next-hop dampening disable

 Example:

RP/0/RP0/CPU0:router(config-rib) # address-family ipv4 next-hop dampening disable
Disables next-hop dampening for IPv4 address families.

Step 3 Use the **commit** or **end** command.

commit —Saves the configuration changes and remains within the configuration session.

- end —Prompts user to take one of these actions:
 - Yes Saves configuration changes and exits the configuration session.
 - No —Exits the configuration session without committing the configuration changes.
 - Cancel Remains in the configuration session, without committing the configuration changes.

Output of show route next-hop Command: Example

The following is sample output from the show route resolving-next-hop command:

```
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/
/CPU0/0
Route metric is 0
```

Enable RCC and LCC On-demand Scan

Perform this task to trigger route consistency checker (RCC) and Label Consistency Checker (LCC) on-demand scan. The on-demand scan can be run on a particular address family (AFI), sub address family (SAFI), table and prefix, vrf, or all prefixes in the table.

SUMMARY STEPS

- 1. Use one of these commands.
 - show rcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name]

- show lcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name]
- **2.** Use one of these commands.
 - clear rcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] log
 - clear lcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] log

DETAILED STEPS

Step 1

Step 2

Use one of these commands. • show rcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] • show lcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] Example: RP/0/RP0/CPU0:router#show rcc ipv6 unicast 2001:DB8::/32 vrf vrf_1 Or RP/0/RP0/CPU0:router#show lcc ipv6 unicast 2001:DB8::/32 vrf vrf_1 Runs on-demand Route Consistency Checker (RCC) or Label Consistency Checker (LCC). Use one of these commands. • clear rcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] log • clear lcc {ipv4 | ipv6} unicast [all] [prefix/mask] [vrf vrf-name] log

Example:

RP/0/RP0/CPU0:router#clear rcc ipv6 unicast log

Or

RP/0/RP0/CPU0:router#show lcc ipv6 unicast log

Clears the log of previous scans.

Enable RCC and LCC Background Scan

Perform this task to run a background scan for Route Consistency Checker (RCC) and Label Consistency Checker (LCC).

SUMMARY STEPS

- 1. configure
- **2.** Use one of these commands:
 - rcc {ipv4 | ipv6} unicast {enable | period milliseconds}
 - lcc {ipv4 | ipv6} unicast {enable | period milliseconds}

- **3.** Use the **commit** or **end** command.
- **4.** Use one of these commands.
 - show rcc {ipv4| ipv6} unicast [summary | scan-id scan-id-value]
 - show lcc {ipv4| ipv6} unicast [summary | scan-id scan-id-value]

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	Use one of these commands:
	 rcc {ipv4 ipv6} unicast {enable period milliseconds} lcc {ipv4 ipv6} unicast {enable period milliseconds}
	Example:
	RP/0/RP0/CPU0:router(config)#rcc ipv6 unicast enable
	RP/0/RP0/CPU0:router(config)#rcc ipv6 unicast period 500
	Or
	RP/0/RP0/CPU0:router(config)#lcc ipv6 unicast enable
	RP/0/RP0/CPU0:router(config)#lcc ipv6 unicast period 500
	Triggers RCC or LCC background scan. Use the period option to control how often the verification be triggered. Each time the scan is triggered, verification is resumed from where it was left out and one buffer's worth of routes or labels are sent to the forwarding information base (FIB).
Step 3	Use the commit or end command.
	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.
	• No —Exits the configuration session without committing the configuration changes.
	• Cancel — Remains in the configuration session, without committing the configuration changes.
Step 4	Use one of these commands.
	 show rcc {ipv4 ipv6} unicast [summary scan-id scan-id-value] show lcc {ipv4 ipv6} unicast [summary scan-id scan-id-value]
	Example:

RP/0/RP0/CPU0:router#show rcc ipv6 unicast statistics scan-id 120

Or

RP/0/RP0/CPU0:router#show lcc ipv6 unicast statistics scan-id 120

Displays statistics about background scans.

- summary—Displays the current ongoing scan id and a summary of the previous few scans.
- scan-id scan-id-value—Displays details about a specific scan.

Enabling RCC and LCC: Example

This example shows how to enable Route Consistency Checker (RCC) background scan with a period of 500 milliseconds between buffers in scans for IPv6 unicast tables:

rcc ipv6 unicast period 500

This example shows how to enable Label Consistency Checker (LCC) background scan with a period of 500 milliseconds between buffers in scans for IPv6 unicast tables:

lcc ipv6 unicast period 500

This example shows how to run Route Consistency Checker (RCC) on-demand scan for subnet 10.10.0.0/16 in vrf1:

show rcc ipv4 unicast 10.10.0.0/16 vrf vrf 1

This example shows how to run Label Consistency Checker (LCC) on-demand scan on all labels for IPv6 prefixes:

show lcc ipv6 unicast all

References for RIB

This section provides additional conceptual information on RIB. It includes the following topics:

- RIB Data Structures in BGP and Other Protocols, on page 102
- RIB Administrative Distance, on page 103
- RIB Statistics, on page 103
- RIB Quarantining, on page 104
- Route and Label Consistency Checker, on page 104

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. 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). RIB processes are not responsible for the BRIB, which are handled by BGP.

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.

This table lists default administrative distances for the common protocols.

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

Table 5: Default Administrative Distances

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 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.

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.

Route and Label Consistency Checker

The Route Consistency Checker and Label Consistency Checker (RCC/LCC) are command-line tools that can be used to verify consistency between control plane and data plane route and label programming in IOS XR software.

Routers in production networks may end up in a state where the forwarding information does not match the control plane information. Possible causes of this include fabric or transport failures between the Route Processor (RP) and the line cards (LCs), or issues with the Forwarding Information Base (FIB). RCC/LCC can be used to identify and provide detailed information about resultant inconsistencies between the control plane and data plane. This information can be used to further investigate and diagnose the cause of forwarding problems and traffic loss.

RCC/LCC can be run in two modes. It can be triggered from using the appropriate command modes as an on-demand, one-time scan (On-demand Scan), or be configured to run at defined intervals in the background during normal router operation (Background Scan). RCC compares the Routing Information Base (RIB) against the Forwarding Information Base (FIB) while LCC compares the Label Switching Database (LSD) against the FIB. When an inconsistency is detected, RCC/LCC output will identify the specific route or label and identify the type of inconsistency detected as well as provide additional data that will assist with further troubleshooting.

RCC runs on the Route Processor. FIB checks for errors on the line card and forwards first the 20 error reports to RCC. RCC receives error reports from all nodes, summarizes them (checks for exact match), and adds it to two queues, soft or hard. Each queue has a limit of 1000 error reports and there is no prioritization in the queue. RCC/LCC logs the same errors (exact match) from different nodes as one error. RCC/LCC compares the errors based on prefix/label, version number, type of error, etc.

On-demand Scan

In On-demand Scan, user requests scan through the command line interface on a particular prefix in a particular table or all the prefixes in the table. The scan is run immediately and the results are published right away. LCC performs on-demand scan on the LSD, where as RCC performs it per VRF.

Background Scan

In Background Scan, user configures the scan that is then left to run in the background. The configuration consists of the time period for the periodic scan. This scan can be configured on either a single table or multiple tables. LCC performs background scan on the LSD, where as RCC performs it either for default or other VRFs.



Implementing RIP

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 to implement basic RIP routing. 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 RFC 2453.

Feature History for Implementing RIP

Release	Modification
Release 6.0.1	This feature was introduced.

- Prerequisites for Implementing RIP, on page 107
- Information About Implementing RIP, on page 107
- Authentication Using Keychain in RIP, on page 111
- How to Implement RIP, on page 112
- Configuration Examples for Implementing RIP, on page 123

Prerequisites for Implementing RIP

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Information About Implementing RIP

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

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.



Note VRF does not allow configuration of a VRF group applied directly under router RIP. A VRF group can be configured if it is applied globally or under VRF.

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 Customizing RIP, on page 114 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).
- 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.

This table lists the 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
External BGP	20
OSPF	110
IS-IS	115
RIP version 1 and 2	120
Internal BGP	200
Unknown	255

Table 6: Default Ad	Iministrative	Distances of	Routing	Protocols
---------------------	---------------	--------------	---------	-----------

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, extended community set, or prefix set. At least 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.

Authentication Using Keychain in RIP

Authentication using keychain in Cisco IOS XR Routing Information Protocol (RIP) provides mechanism to authenticate all RIP protocol traffic on RIP interface, based keychain authentication. This mechanism uses the Cisco IOS XR security keychain infrastructure to store and retrieve secret keys and use it to authenticate in-bound and out-going traffic on per-interface basis.

Keychain management is a common method of authentication to configure shared secrets on all entities that exchange secrets such as keys, before establishing trust with each other. Routing protocols and network management applications on Cisco IOS XR software often use authentication to enhance security while communicating with peers.



- The Cisco IOS XR software system security component implements various system security features including keychain management. Refer these documents for detailed information on keychain management concepts, configuration tasks, examples, and command used to configure keychain management.
 - Implementing Keychain Management module in System Security Configuration Guide for Cisco NCS 5000 Series Routers
 - Keychain Management Commands module in System Security Command Reference for Cisco NCS 5000
 Series Routers



Note

The keychain by itself has no relevance; therefore, it must be used by an application that needs to communicate by using the keys (for authentication) with its peers. The keychain provides a secure mechanism to handle the keys and rollover based on the lifetime. The Cisco IOS XR keychain infrastructure takes care of the hit-less rollover of the secret keys in the keychain.

Once you have configured a keychain in the IOS XR keychain database and if the same has been configured on a particular RIP interface, it will be used for authenticating all incoming and outgoing RIP traffic on that interface. Unless an authentication keychain is configured on a RIP interface (on the default VRF or a non-default VRF), all RIP traffic will be assumed to be authentic and authentication mechanisms for in-bound RIP traffic and out-bound RIP traffic will be not be employed to secure it.

RIP employs two modes of authentication: keyed message digest mode and clear text mode. Use the **authentication keychain** *keychain-name* **mode** {**md5** | **text**} command to configure authentication using the keychain mechanism.

In cases where a keychain has been configured on RIP interface but the keychain is actually not configured in the keychain database or keychain is not configured with MD5 cryptographic algorithm, all incoming RIP packets on the interface will be dropped. Outgoing packets will be sent without any authentication data.

In-bound RIP Traffic on an Interface

These are the verification criteria for all in-bound RIP packets on a RIP interface when the interface is configured with a keychain.

If	Then	
The keychain configured on the RIP interface does not exist in the keychain database	The packet is dropped. A RIP component-level debug message is be logged to provide the specific details of the authentication failure.	
The keychain is not configured with a MD5 cryptographic algorithm	The packet is dropped. A RIP component-level debug message is be logged to provide the specific details of the authentication failure.	
The Address Family Identifier of the first (and only the first) entry in the message is not 0xFFFF, then authentication is not in use	The packet will be dropped. A RIP component-level debug message is be logged to provide the specific details of the authentication failure.	
The MD5 digest in the 'Authentication Data' is found to be invalid	The packet is dropped. A RIP component-level debug message is be logged to provide the specific details of the authentication failure.	
Else, the packet is forwarded for the rest of the processing.		

Out-bound RIP Traffic on an Interface

These are the verification criteria for all out-bound RIP packets on a RIP interface when the interface is configured with a keychain.

If	Then
The keychain configured on the RIP interface exists in the keychain database	The RIP packet passes authentication check at the remote/peer end, provided the remote router is also configured to authenticate the packets using the same keychain.
The keychain is configured with a MD5 cryptographic algorithm	The RIP packet passes authentication check at the remote/peer end, provided the remote router is also configured to authenticate the packets using the same keychain.
Else, RIP packets fail authentication check.	

How to Implement RIP

This section contains instructions for the following tasks:



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.

Before you begin

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 interface-path-id
- 6. receive version $\{1 \mid 2 \mid 12\}$
- 7. send version $\{1 \mid 2 \mid 12\}$
- 8. Use the commit or end command.

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	router rip	Configures a RIP routing process.
	Example:	
	RP/0/RP0/CPU0:router(config)# router rip	
Step 3	neighbor ip-address	(Optional) Defines a neighboring router with which to
	Example:	exchange RIP protocol information.
	RP/0/RP0/CPU0:router(config-rip)# neighbor 172.160.1.2	
Step 4	broadcast-for-v2	(Optional) Configures RIP to send only Version 2 packets
	Example:	to the broadcast IP address. This command can be applied at the interface or level.
	RP/0/RP0/CPU0:router(config-rip)# broadcast-for-v2	
Step 5	interface type interface-path-id	(Optional) Defines the interfaces on which the RIP routing
	Example:	protocol runs.
	RP/0/RP0/CPU0:router(config-rip)# interface HundredGigE 0/1/0/3	

	Command or Action	Purpose
Step 6	receive version { 1 2 1 2 }	(Optional) Configures an interface to accept packets that
	Example:	are:
		• Only RIP v1
	<pre>RP/0/RP0/CPU0:router(config-rip-if)# receive version 1 2</pre>	• Only RIP v2
		• Both RIP v1 and RIP v2
Step 7	send version { 1 2 1 2 }	(Optional) Configures an interface to send packets that are
	Example:	• Only RIP v1
	<pre>RP/0/RP0/CPU0:router(config-rip-if)# send version 1 2</pre>	• Only RIP v2
		• Both RIP v1 and RIP v2
Step 8	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Customizing 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 interface-path-id
- 8. metric-zero-accept
- 9. split-horizon disable
- **10.** poison-reverse
- **11.** Use the **commit** or **end** command.

DETAILED STEPS

	Command or Action	Purpose	
Step 1	configure	Enters mode.	
	Example:		
	RP/0/RP0/CPU0:router# configure		
Step 2	router rip	Configures a RIP routing process.	
	Example:		
	RP/0/RP0/CPU0:router(config)# router rip		
Step 3	auto-summary	(Optional) Enables automatic route summarization of	
	Example:	subnet routes into network-level routes.	
	<pre>RP/0/RP0/CPU0:router(config-rip)# auto-summary</pre>	• By default, auto-summary is disabled.	
		Note If you have disconnected subnets, use the no keyword to disable automatic route summarization and permit software to send subnet and host routing information across classful network boundaries.	
Step 4	timers basic update invalid holddown flush	(Optional) Adjusts RIP network timers.	
	Example:	Note To view the current and default timer values, view output from the show rip command.	
	<pre>RP/0/RP0/CPU0:router(config-rip)# timers basic 5 15 15 30</pre>		
Step 5	output-delay delay	(Optional) Changes the interpacket delay for the RIP	
	Example:	updates sent.	
	<pre>RP/0/RP0/CPU0:router(config-rip)# output-delay 10</pre>	Note 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	nsf	(Optional) ensures continuous forwarding even after a R	
	Example:	process is shutdown or restart.	
	RP/0/RP0/CPU0:router(config-rip)# nsf		
Step 7	interface type interface-path-id	(Optional) Defines the interfaces on which the RIP routi	
	Example:	protocol runs.	
	RP/0/RP0/CPU0:router(config-rip)# interface HundredGigE 0/1/0/3		

Command or Action	Purpose
metric-zero-accept Example:	(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).
<pre>RP/0/RP0/CPU0:router(config-rip-if)# metric-zero-accept</pre>	
split-horizon disable	(Optional) Disables the split horizon mechanism.
Example:	• By default, split horizon is enabled.
<pre>RP/0/RP0/CPU0:router(config-rip-if)# split-horizon disable</pre>	• In general, we do not recommend changing the state of the default for the split-horizon command, unless you are certain that your application requires a change to properly advertise routes.
poison-reverse	Enables poison reverse processing of RIP router updates.
Example:	
RP/0/RP0/CPU0:router(config-rip-if)# poison-reverse	
Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.
	• No —Exits the configuration session without committing the configuration changes.
	• Cancel —Remains in the configuration session, without committing the configuration changes.
-	<pre>metric-zero-accept Example: RP/0/RP0/CPU0:router(config-rip-if)# metric-zero-accept split-horizon disable Example: RP/0/RP0/CPU0:router(config-rip-if)# split-horizon disable poison-reverse Example: RP/0/RP0/CPU0:router(config-rip-if)# poison-reverse</pre>

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 interface-path-id
- 5. passive-interface
- **6.** exit
- 7. interface type interface-path-id
- **8.** route-policy $\{ in \mid out \}$
- **9.** Use the **commit** or **end** command.

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	router rip	Configures a RIP routing process.
	Example:	
	RP/0/RP0/CPU0:router(config)# router rip	
Step 3	neighbor ip-address	(Optional) Defines a neighboring router with which to
	Example:	exchange RIP protocol information.
	<pre>RP/0/RP0/CPU0:router(config-rip)# neighbor 172.160.1.2</pre>	
Step 4	interface type interface-path-id	(Optional) Defines the interfaces on which the RIP routin protocol runs.
	Example:	
	RP/0/RP0/CPU0:router(config-rip)# interface HundredGigE 0/1/0/3	
Step 5	passive-interface	(Optional) Suppresses the sending of RIP updates on an interface, but not to explicitly configured neighbors.
	Example:	
	<pre>RP/0/RP0/CPU0:router(config-rip-if)# passive-interface</pre>	
Step 6	exit	(Optional) Returns the router to the next higher
	Example:	configuration mode.

	Command or Action	Purpose
	RP/0/ /CPU0:router(config-rip-if)# exit	
Step 7	interface type interface-path-id Example:	(Optional) Defines the interfaces on which the RIP routing protocol runs.
	<pre>RP/0/RP0/CPU0:router(config-rip)# interface HundredGigE 0/1/0/4</pre>	
Step 8	<pre>route-policy { in out } Example:</pre>	(Optional) Applies a routing policy to updates advertised to or received from a RIP neighbor.
	<pre>RP/0/RP0/CPU0:router(config-rip-if)# route-policy out</pre>	
Step 9	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

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

- 1. configure
- 2. route-policy name
- 3. set rip-metric number
- 4. end-policy
- 5. Use the commit or end command.

- 6. configure
- 7. router rip
- **8.** route-policy *route-policy-name* { **in** | **out** }
- **9.** Use the **commit** or **end** command.

DETAILED STEPS

I

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	route-policy name	Defines a route policy and enters route-policy configuration mode.
	Example:	
	RP/0/RP0/CPU0:router(config)# route-policy IN-IPv4	
Step 3	set rip-metric number	(Optional) Sets the RIP metric attribute.
	Example:	
	RP/0/RP0/CPU0:router(config-rpl)# set rip metric 42	
Step 4	end-policy	Ends the definition of a route policy and exits route-policy
	Example:	configuration mode.
	RP/0/RP0/CPU0:router(config-rpl)# end-policy	
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.
Step 6	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 7	router rip	Configures a RIP routing process.
	Example:	

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config)# router rip	
Step 8	route-policy route-policy-name { in out } Example:	Applies a routing policy to updates advertised to or received from an RIP neighbor.
	<pre>RP/0/RP0/CPU0:router(config-rip)# route-policy rp1 in</pre>	
Step 9	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Configuring RIP Authentication Keychain

Configuring RIP Authentication Keychain for IPv4 Interface on a Non-default VRF

Perform this task to configure a RIP authentication keychain for IPv4 interface on a non-default VRF.

Before you begin

All keychains need to be configured in Cisco IOS XR keychain database using configuration commands described in *Implementing Keychain Management* module of *System Security Configuration Guide for Cisco NCS 5000 Series Routers* before they can be applied to a RIP interface/VRF.

The **authentication keychain** *keychain-name* and **mode md5** configurations will accept the name of a keychain that has not been configured yet in the IOS XR keychain database or a keychain that has been configured in IOS XR keychain database without MD5 cryptographic algorithm. However, in both these cases, all incoming packets on the interface will be dropped and outgoing packets will be sent without authentication data.

SUMMARY STEPS

- 1. configure
- 2. router rip
- **3. vrf**_*name*
- 4. interface type interface-path-id
- **5.** Use one of these commands:
 - authentication keychain keychain-name mode md5
 - authentication keychain keychain-name mode text

6. Use the **commit** or **end** command.

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	router rip	Configures a RIP routing process.
	Example:	
	RP/0/RP0/CPU0:router(config)#router rip	
Step 3	vrf vrf_name	Configures a non-default VRF
	Example:	
	RP/0/RP0/CPU0:router(config-rip)#vrf vrf_rip_auth	
Step 4	interface type interface-path-id	Defines the interface on which the RIP routing protocol
-	Example:	runs.
	RP/0/RP0/CPU0:router(config-rip-vrf)#interface HundredGigE 0/1/0/3	
Step 5	Use one of these commands:	Configures an authentication keychain mode for RIP.
	 • authentication keychain keychain-name mode md5 • authentication keychain keychain-name mode text 	• md5—Keyed message digest (md5) authentication mode
	Example:	• text—Clear text authentication mode
	<pre>RP/0/RP0/CPU0:router(config-rip-if)#authentication keychain key1 mode md5</pre>	
	Or	
	<pre>RP/0/RP0/CPU0:router(config-rip-if)#authentication keychain key1 mode text</pre>	
Step 6	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel — Remains in the configuration session, without committing the configuration changes.

Configuring RIP Authentication Keychain for IPv4 Interface on Default VRF

Perform this task to configure a RIP authentication keychain for IPv4 interface (on the default VRF).

Before you begin

All keychains need to be configured in Cisco IOS XR keychain database using configuration commands described in *Implementing Keychain Management* module of *System Security Configuration Guide for Cisco NCS 5000 Series Routers* before they can be applied to a RIP interface/VRF.

The **authentication keychain** *keychain-name* and **mode md5** configurations will accept the name of a keychain that has not been configured yet in the IOS XR keychain database or a keychain that has been configured in IOS XR keychain database without MD5 cryptographic algorithm. However, in both these cases, all incoming packets on the interface will be dropped and outgoing packets will be sent without authentication data.

SUMMARY STEPS

- 1. configure
- 2. router rip
- **3.** interface type interface-path-id
- **4.** Use one of these commands:
 - authentication keychain keychain-name mode md5
 - authentication keychain keychain-name mode text
- 5. Use the commit or end command.

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	router rip	Configures a RIP routing process.
	Example:	
	RP/0/RP0/CPU0:router(config) #router rip	
Step 3	interface type interface-path-id	Defines the interface on which the RIP routing protoco runs.
	Example:	
	RP/0/RP0/CPU0:router(config-rip)#interface HundredGigE 0/1/0/3	
Step 4	Use one of these commands:	Configures an authentication keychain mode for RIP.
	 authentication keychain keychain-name mode md5 authentication keychain keychain-name mode text 	
	Example:	• text—Clear text authentication mode

	Command or Action	Purpose
	<pre>RP/0/RP0/CPU0:router(config-rip-if)#authentication keychain key1 mode md5</pre>	
	Or	
	<pre>RP/0/RP0/CPU0:router(config-rip-if)#authentication keychain keyl mode text</pre>	
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		 No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Configuration Examples for Implementing RIP

This section provides the following configuration examples:

Configuring a Basic RIP Configuration: Example

The following example shows two Gigabit Ethernet interfaces configured with RIP.

```
interface TenGigE 0/3/0/0
ipv4 address 172.16.0.1 255.255.255.0
!
interface TenGigE 0/3/0/1
ipv4 address 172.16.2.12 255.255.255.0
!
router rip
interface TenGigE 0/3/0/0
!
interface TenGigE 0/3/0/1
!
```

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 HundredGigE 0/1/0/3
```

```
!
vrf vpn0
interface HundredGigE 0/1/0/4
!
!
vrf vpn1
interface HundredGigE 0/1/0/5
!
!
```

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 HundredGigE 0/1/0/3
!
vrf vpn0
interface HundredGigE 0/1/0/4
!
timers basic 10 30 30 45
!
vrf vpn1
interface HundredGigE 0/1/0/5
!
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, 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 HundredGigE 0/1/0/3
!
redistribute connected
default-metric 3
!
vrf vpn1
interface HundredGigE 0/1/0/4
!
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 pfl
 10.1.0.0/24
end-set
1
prefix-set pf2
 150.10.1.0/24
end-set
1
route-policy policy in
 if destination in pf1 then
   pass
 endif
end-policy
1
route-policy pass-all
 pass
end-policy
!
route-policy infil
 if destination in pf2 then
   add rip-metric 2
   pass
 endif
end-policy
1
router rip
interface HundredGigE 0/1/0/3
 route-policy policy in in
 1
interface HundredGigE 0/1/0/4
!
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 HundredGigE 0/1/0/3
passive-interface
!
interface HundredGigE 0/1/0/4
!
neighbor 172.17.0.1
neighbor 172.18.0.5
!
```



Implementing Routing Policy

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.

This module describes how routing protocols make decisions to advertise, aggregate, discard, distribute, export, hold, import, redistribute and modify the routes based on configured routing policy.

The routing policy language (RPL) provides 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 streamlines the routing policy configuration, reduces system resources required to store and process these configurations, and simplifies troubleshooting.

- Restrictions for Implementing Routing Policy, on page 127
- Define Route Policy, on page 128
- Attach Routing Policy to BGP Neighbor, on page 130
- Modify Routing Policy Using Text Editor, on page 131
- References for Routing Policy, on page 134

Restrictions for Implementing Routing Policy

These restrictions apply when working with Routing Policy Language implementation:

- Border Gateway Protocol (BGP), integrated Intermediate System-to-Intermediate System (IS-IS), or Open Shortest Path First (OSPF) must be configured in your network.
- An individual policy definition of up to 1000 statements are supported. The total number of statements within a policy can be extended to 4000 statements using hierarchical policy constructs. However, this limit is restricted with the use of **apply** statements.
- When a policy that is attached directly or indirectly to an attach point needs to be modified, a single **commit** operation cannot be performed when:
 - Removing a set or policy referred by another policy that is attached to any attach point directly or indirectly.
 - Modifying the policy to remove the reference to the same set or policy that is getting removed.

The **commit** must be performed in two steps:

- 1. Modify the policy to remove the reference to the policy or set and then commit.
- 2. Remove the policy or set and commit.
- Per-vrf label mode is not supported for Carrier Supporting Carrier (CSC) network with internal and external BGP multipath setup.
- You cannot change the next hop address to an IPv6 address through RPL policy for a route that starts from an IPv4 peer.

Define 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.
- Modifying the RPL scale configuration may take a long time.
- BGP may crash either due to large scale RPL configuration changes, or during consecutive RPL changes. To avoid BGP crash, wait until there are no messages in the BGP In/Out queue before committing further changes.

Tip You can programmatically configure the route policy using <code>openconfig-routing-policy.yang</code> OpenConfig data model. To get started with using data models, see the *Programmability* Configuration Guide.

SUMMARY STEPS

- 1. configure
- **2.** route-policy name [parameter1, parameter2,..., parameterN]
- **3**. end-policy
- 4. Use the commit or end command.

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 2 route-policy name [parameter1, parameter2,..., parameterN]

Example:

RP/0/RP0/CPU0:router(config) # route-policy sample1

Enters route-policy configuration mode.

After the route-policy has been entered, a group of commands can be entered to define the route-policy.

Step 3 end-policy

Example:

RP/0/RP0/CPU0:router(config-rpl)# end-policy

Ends the definition of a route policy and exits route-policy configuration mode.

Step 4 Use the **commit** or **end** command.

commit —Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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
```

Attach Routing Policy to BGP Neighbor

This task explains how to attach a routing policy to a BGP neighbor.

Before you begin

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

- 1. configure
- 2. router bgp as-number
- 3. neighbor ip-address
- 4. address-family { ipv4 unicast || ipv6 unicast |} address-family { ipv4 | ipv6 } unicast
- **5.** route-policy *policy-name* { in | out }
- 6. Use the commit or end command.

DETAILED STEPS

Step 1 configure Example:

RP/0/RP0/CPU0:router# configure

Enters mode.

Step 2 router bgp *as-number*

Example:

RP/0/RP0/CPU0:router(config) # router bgp 125

Configures a BGP routing process and enters router configuration mode.

• The *as-number* 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.

Step 3 neighbor *ip-address*

Example:

RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.20

Specifies a neighbor IP address.

Step 4 address-family { ipv4 unicast || ipv6 unicast |} address-family { ipv4 | ipv6 } unicast

Example:

RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast

Specifies the address family.

Step 5route-policy policy-name{ in | out }

Example:

RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy example1 in

Attaches the route-policy, which must be well formed and predefined.

Step 6 Use the commit or end command.

commit —Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

Modify Routing Policy Using Text Editor

This task explains how to modify an existing routing policy using a text editor.

SUMMARY STEPS

- 1. 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]
- **2**. show rpl route-policy [*name* [detail] | states | brief]
- **3.** show rpl prefix-set [name | states | brief]

DETAILED STEPS

 Step 1
 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]

Example:

RP/0/RP0/CPU0:router# edit route-policy sample1

Identifies the route policy, prefix set, AS path set, community set, or extended community set name to be modified.

- 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.
- After editing with Nano, save the editor buffer and exit the editor by using the Ctrl-X keystroke.
- 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.

• 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.

 Step 2
 show rpl route-policy
 [name [detail] | states | brief]

Example:

RP/0/RP0/CPU0:router# show rpl route-policy sample2

(Optional) Displays the configuration of a specific named route policy.

- Use the **detail** keyword to display all policies and sets that a policy uses.
- Use the states keyword to display all unused, inactive, and active states.
- Use the **brief** keyword to list the names of all extended community sets without their configurations.

Step 3 show rpl prefix-set [name | states | brief]

Example:

RP/0/RP0/CPU0:router# show rpl prefix-set prefixset1

(Optional) Displays the contents of a named prefix set.

• To display the contents of a named AS path set, community set, or extended community set, replace the **prefix-set** keyword with **as-path-set**, **community-set**, or **extcommunity-set**, respectively.

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 bqp 2
 neighbor 10.0.1.2 address-family ipv4 unicast route-policy inbound-tx in
```

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.

```
prefix-set bogon
  10.0.0/8 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
Т
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 iqp
 set community (2:333)
end-policy
route-policy set-lpref-prepend($lpref,$as,$prependcnt)
 set local-preference $lpref
 prepend as-path $as $prependent
end-policy
```

References for Routing Policy

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

Routing Policy Language

This section contains the following information:

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.

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, 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.



Note Regular expression matching is CPU intensive. The policy performance can be substantially improved by either collapsing the regular expression patterns together to reduce the total number of regular expression invocations or by using equivalent native as-path match operations such as 'as-path neighbor-is', 'as-path originates-from' or 'as-path passes-through'.

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.

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.

Wildcards (*) and regular expressions are allowed for extended community set elements.

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 RT

An rt set is an extcommunity set used to store BGP Route Target (RT) 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)
  ($ipadrr:666, 1234:$tag, 1.2.3.4:777, $tag2:777)
```

These options are supported under extended community set RT:

```
RP/0/RP0/CPU0:router(config) #extcommunity-set rt rt set
RP/0/RP0/CPU0:router(config-ext)#?
 #-remark Remark beginning with '#'
                Wildcard (any community or part thereof)
 <1-4294967295> 32-bit decimal number
 <1-65535> 16-bit decimal number
 A.B.C.D/M:N Extended community - IPv4 prefix format
 A.B.C.D:N Extended community - IPv4 format
 ASN:N
                Extended community - ASPLAIN format
               Extended community - ASDOT format
 X.Y:N
               Discard RPL definition and return to top level config
 abort.
 dfa-regex
               DFA style regular expression
 end-set
               End of set definition
 exit
                Exit from this submode
                Traditional IOS style regular expression
  ios-regex
  show
                Show partial RPL configuration
```

L

Option	Description
#-remark	Remark beginning with '#'
*	Wildcard (any community or part thereof)
<1-4294967295>	32-bit decimal number
<1-65535>	16-bit decimal number
A.B.C.D/M:N	Extended community - IPv4 prefix format
A.B.C.D:N	Extended community - IPv4 format
ASN:N	Extended community - ASPLAIN format
X.Y:N	Extended community - ASDOT format
abort	Discard RPL definition and return to top level config
dfa-regex	DFA style regular expression
end-set	End of set definition
exit	Exit from this submode
ios-regex	Traditional IOS style regular expression
show	Show partial RPL configuration

Named Form for Extcommunity-set Soo

A soo set is an extcommunity set used to store BGP Site-of-Origin (SoO) extended community type communities:

extcommunity-set soo a_soo_set
1.1.1:100,
 100:200
end-set

These options are supported under extended community set Soo:

```
RP/0/RP0/CPU0:router(config)#extcommunity-set soo soo_set
RP/0/RP0/CPU0:router(config-ext)#?
#-remark Remark beginning with '#'
* Wildcard (any community or part thereof)
<1-4294967295> 32-bit decimal number
<1-65535> 16-bit decimal number
A.B.C.D/M:N Extended community - IPv4 prefix format
A.B.C.D:N Extended community - IPv4 format
ASN:N Extended community - ASPLAIN format
X.Y:N Extended community - ASPLAIN format
abort Discard RPL definition and return to top level config
dfa-regex DFA style regular expression
end-set End of set definition
exit Exit from this submode
ios-regex Traditional IOS style regular expression
show Show partial RPL configuration
```

Option	Description
#-remark	Remark beginning with '#'
*	Wildcard (any community or part thereof)
<1-4294967295>	32-bit decimal number
<1-65535>	16-bit decimal number
A.B.C.D/M:N	Extended community - IPv4 prefix format
A.B.C.D:N	Extended community - IPv4 format
ASN:N	Extended community - ASPLAIN format
X.Y:N	Extended community - ASDOT format
abort	Discard RPL definition and return to top level config
dfa-regex	DFA style regular expression
end-set	End of set definition
exit	Exit from this submode
ios-regex	Traditional IOS style regular expression
show	Show partial RPL configuration

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 **g**reater than or **e**qual 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 **l**ess than or **e**qual 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 **e**qual 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.

ACL Support in RPL Prefix Sets

Access Control List (ACL) type prefix set entries holds IPv4 or IPv6 prefix match specifications, each of which has an address and a wildcard mask. The address and wildcard mask is a standard dotted-decimal IPv4 or colon-separated hexadecimal IPv6 address. The set of bits to be matched are provided in the form of wildcard also called as inverted mask in which a binary 0 means a mandatory match and binary 1 means a do not match condition. The prefix set allows to specify contiguous and non-contiguous set of bits that should be matched in any route.

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.

Pass PolicyPass Policy

The following example shows how the policy accepts all presented routes without modifying the routes.

```
route-policy quickstart-pass
pass
end-policy
```

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
```

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, policies can also be defined that allow for parameterization the attributes at attach points. Parameterization is supported at all attach points.

In the following example, 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 gbpathtype 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

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 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:

```
ifmed eq 12 then
set med 42
if med eq 42 then
drop
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.

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 10.0.0/8 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.

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 \text{ 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)
```

lf

```
end-policy
```

lf

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

This section describes each of the BGP policy attach points and provides a summary of the BGP attributes and operators.

Additional-Path

The additional-path attach point provides increased control based on various attribute match operations. This attach point is used to decide whether a route-policy should be used to select additional-paths for a BGP speaker to be able to send multiple paths for the prefix.

The add path enables BGP prefix independent convergence (PIC) at the edge routers.

This example shows how to set a route-policy "add-path-policy" to be used for enabling selection of additional paths:

```
router bgp 100
address-family ipv4 unicast
additional-paths selection route-policy add-path-policy
```

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 are dampened than the routes that are less specific than /25.



Note When the dampening policy runs for a route, then the last "set dampening" statement that is encountered, takes effect.

- If a "drop" statement is encountered, then the route is not dampened; even if the "set dampening" statement is encountered.
- If a "pass" or "done" statement is encountered but not the "set dampening" statement, then the route is dampened using the default dampening parameters.

For example:

- When policy1 applies another policy that is called policy2 and if a "pass" statement is encountered in policy2, then policy2 exits and continues to execute policy1.
- If a "done" statement is encountered in policy2, then both policy1 and policy2 exits immediately.

```
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:

```
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
!
```

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
*> 10.0.0.0 10.0.1.2
                                    Metric LocPrf Weight Path
                                                                 03?
                                              10

        *> 10.0.0/9
        10.0.1.2

        *> 10.0.0.0/10
        10.0.1.2

        *> 10.0.0.0/11
        10.1.2

                                                 10
                                                                  03?
                                                 10
                                                                 03?
                                                 10
                                                                 03?
*> 10.1.0.0/16
                       10.0.1.2
                                                 10
                                                                 03?
                                                10
*> 10.3.30.0/24
                     10.0.1.2
                                                                 03?
*> 10.3.30.128/25 10.0.1.2
                                               10
                                                                 03?
                   10.0.1.2
                                                10
*> 10.128.0.0/9
                                                                 03?
                                                       555 0 100 e
                      10.0.101.2
                                              1000 555
1000 555
*> 10.255.0.0/24
*> 10.255.64.0/24
                       10.0.101.2
. . . .
```

....

show bgp policy route-policy test2

10.0.0.0/8 is advertised to 10.0.101.2

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 (172.16.1.0/24, 172.16.5.0/24, 172.16.11.0/24). 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
  172.16.1.0/24,
  172.16.5.0/24.
 172.16.11.0/24
end-set
route-policy policy-orf
  if orf prefix in orf-preset then
    drop
  endif
  if orf prefix in (172.16.3.0/24, 172.16.7.0/24, 172.16.13.0/24) 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 the BGP next-hop tracking feature using a route policy to monitor static or connected routes with the prefix 10.0.0.0 and prefix length 8.

```
route-policy nxthp_policy_A
    if destination in (10.0.0.0/8) 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 with prefix length 8; therefore, the debug output shows up only for that prefix.

```
route-policy policy_b
  if destination in (10.0.0.0/8) then
   pass
  else
    drop
  endif
end-policy
```

debug bgp update policy_b

BGP Attributes and Operators

This table summarizes the BGP attributes and operators per attach points.

Table 7: BGP Attributes and Operators

Attach Point	Attribute	Match	Set
aggregation	as-path	in	—
		is-local	
		length	
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	—
	as-path-unique-length	is, ge, le, eq	—
	community	is-empty	set
		matches-any	set additive
		matches-every	delete in
			delete not in
			delete all
	destination	in	—
	extcommunity cost	—	set
			set additive
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	setset +set -
	next-hop	in	set
	origin	is	set
	source	in	—
	suppress-route	—	suppress-route
	weight	_	set

I

Attach Point	Attribute	Match	Set
allocate-label	as-path	in	—
		is-local	
		length	
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	—
	as-path-unique-length	is, ge, le, eq	—
	community	is-empty	—
		matches-any	
		matches-every	
	destination	in	—
	label		set
	local-preference	is, ge, le, eq	—
	med	is, eg, ge, le	—
	next-hop	in	—
	origin	is	—
	source	in	—
clear-policy	as-path	in	—
		is-local	
		length	
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	_
	as-path-unique-length	is, ge, le, eq	

Attach Point	Attribute	Match	Set
dampening	as-path	in	—
		is-local	
		length	
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	—
	as-path-unique-length	is, ge, le, eq	—
	community	is-empty	—
		matches-any	
		matches-every	
	dampening	/	set dampening
	destination	in	—
	local-preference	is, ge, le, eq	—
	med	is, eg, ge, le	—
	next-hop	in	—
	origin	is	—
	source	in	—
debug	destination	in	—
default	med	—	set
originate			set +
			set -
	rib-has-route	in	—

¢

I

Attach Point	Attribute	Match	Set
neighbor-in	as-path	in	prepend
		is-local	prepend most-recent
		length	remove as-path private-as
		NA	replace
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	
	as-path-unique-length	is, ge, le, eq	—
	communitycommunity with 'peeras'	is-empty	set
		matches-any	set additive
		matches-every	delete-in
			delete-not-in
			delete-all
	destination	in	
	extcommunity cost		set
			set additive
	extcommunity rt	is-empty	set
		matches-any	additive
		matches-every	delete-in
		matches-within	delete-not-in
			delete-all
	extcommunity soo	is-empty	-
		matches-any	
		matches-every	
		matches-within	
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	set
			set +
			set -

Attach Point	Attribute	Match	Set
	next-hop	in	set
			set peer address
	origin	is	set
	route-aggregated	route-aggregated	NA
	source	in	—
	weight		set

I

Attach Point	Attribute	Match	Set
neighbor-out	as-path	in	prepend
		is-local	prepend most-recent
		length	remove as-path private-as
		_	replace
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	—
	as-path-unique-length	is, ge, le, eq	—
	communitycommunity with 'peeras'	is-empty	set
		matches-any	set additive
		matches-every	delete-in
			delete-not-in
			delete-all
	destination	in	—
	extcommunity cost	_	set
			set additive
	extcommunity rt	is-empty	set
		matches-any	additive
		matches-every	delete-in
		matches-within	delete-not-in
			delete-all
	extcommunity soo	is-empty	_
		matches-any	
		matches-every	
		matches-within	
	local-preference	is, ge, le, eq	set
	med	is, eg, ge, le	

Attach Point	Attribute	Match	Set
			set
			set +
			set -
			set max-unreachable
			set igp-cost
	next-hop	in	set
			set self
	origin	is	set
	path-type	is	—
	rd	in	—
	route-aggregated	route-aggregated	—
	source	in	—
	unsuppress-route		unsuppress-route
	vpn-distinguisher		set
neighbor-orf	orf-prefix	in	n/a

I

Attach Point	Attribute	Match	Set
network	as-path	—	prepend
	community	—	set
			set additive
			delete-in
			delete-not-in
			delete-all
	destination	in	—
	extcommunity cost	_	set
			set additive
	mpls-label	route-has-label	—
	local-preference	—	set
	med	—	set
			set+
			set-
	next-hop	in	set
	origin	—	set
	route-type	is	—
	tag	is, ge, le, eq	—
	weight	_	set
next-hop	destination	in	—
	protocol	is,in	
	source	in	—

Attach Point	Attribute	Match	Set
redistribute	as-path	—	prepend
	community	—	set
			set additive
			delete in
			delete not in
			delete all
	destination	in	—
	extcommunity cost	—	setset additive
	local-preference	—	set
	med	—	set
			set+
			set-
	next-hop	in	set
	origin	—	set
	mpls-label	route-has-label	—
	route-type	is	—
	tag	is, eq, ge, le	—
	weight	—	set
retain-rt	extcommunity rt	is-empty	—
		matches-any	
		matches-every	
		matches-within	

Attach Point	Attribute	Match	Set
show	as-path	in	
		is-local	
		length	
		neighbor-is	
		originates-from	
		passes-through	
		unique-length	
	as-path-length	is, ge, le, eq	—
	as-path-unique-length	is, ge, le, eq	—
	community	is-empty	—
		matches-any	
		matches-every	
	destination	in	—
	extcommunity rt	is-empty	—
		matches-any	
		matches-every	
		matches-within	
	extcommunity soo	is-empty	
		matches-any	
		matches-every	
		matches-within	
	med	is, eg, ge, le	—
	next-hop	in	—
	origin	is	—
	source	in	

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.

This table summarizes which operations are valid and where they are valid.

Command	import	export	aggregation	redistribution
prepend as-path most-recent	eBGP only	eBGP only	n/a	n/a
replace as-path	eBGP only	eBGP only	n/a	n/a
set med igp-cost	forbidden	eBGP only	forbidden	forbidden
set weight	n/a	forbidden	n/a	n/a
suppress	forbidden	forbidden	n/a	forbidden

Table 8: Restricted BGP Operations by Attach Point

The **bgp cluster-id** option is used to configure the router as one of the route reflectors serving the cluster. The **cluster-id** option is also available in the BGP neighbor address-family (config-bgp-nbr-af) mode. To enable a router to accept BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs, use the **cluster-id allow-equal** command. You must use this command with care to avoid routing loops.

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 10.0.0/8 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
  .
  .
```

OSPF Policy Attach Points

This section describes each of the OSPF policy attach points and provides a summary of the OSPF attributes and operators.

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 10.105.3.0/24, 10.105.7.0/24, 10.105.13.0/24, it is accepted. If the prefix matches any of 10.106.3.0/24, 10.106.13.0/24, it is dropped.

```
route-policy OSPF-area-in
   if destination in (10
.105.3.0/24, 10
.105.7.0/24, 10
.105.13.0/24) then
     drop
    endif
   if destination in (10
.106.3.0/24, 10
.106.7.0/24, 10
.106.13.0/24) then
     pass
    endif
 end-policy
 router ospf 1
   area 1
      route-policy OSPF-area-in in
```

Area-out

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 10.105.3.0/24, 10.105.7.0/24, 10.105.13.0/24, it is announced. If the prefix matches any of 10.105.3.0/24, 10.105.7.0/24, 10.105.13.0/24, it is dropped and not announced.

```
route-policy OSPF-area-out
   if destination in (10
.105.3.0/24, 10
.105.7.0/24, 10
.105.13.0/24) then
     drop
   endif
   if destination in (10
.105.3.0/24, 10
.105.7.0/24, 10
.105.13.0/24) then
     pass
   endif
 end-policy
 router ospf 1
   area 1
     route-policy OSPF-area-out out
```

OSPF Attributes and Operators

This table summarizes the OSPF attributes and operators per attach points.

Attach Point	Attribute	Match	Set
default-information originate	ospf-metric	—	set
originate	metric-type		set
	tag	—	set
	rib-has-route	in	

Table 9: OSPF Attributes an	d Operators
-----------------------------	-------------

Attach Point	Attribute	Match	Set
redistribute	destination	in	—
	metric-type	—	set
	ospf-metric		set
	next-hop	in	—
	mpls-label	route-has-label	—
	rib-metric	is, le, ge, eq	n⁄a
	route-type	is	—
	tag	is, eq, ge, le	set
area-in	destination	in	—
area-out	destination	in	—
spf-prefix-priority	destination	in	n⁄a
	spf-priority	n/a	set
	tag	is, le, ge, eq	n⁄a

Distribute-list in

The distribute-list in attach point within OSPF allows use of route policies to filter OSPF prefixes. The distribute-list in route-policy can be configured at OSPF instance, area, and interface levels. The route-policy used in the distribute-list in command supports match statements, "destination" and "rib-metric". The "set" commands are not supported in the route-policy.

These are examples of valid route-policies for "distribute-list in":

```
route-policy DEST
  if destination in (10.10.10.10/32) then
    drop
  else
   pass
  endif
end-policy
route-policy METRIC
  if rib-metric ge 10 and rib-metric le 19 then
    drop
  else
   pass
  endif
end-policy
prefix-set R-PFX
  10.10.10.30
```

```
end-set
route-policy R-SET
if destination in R-PFX and rib-metric le 20 then
pass
else
drop
endif
end-policy
```

OSPFv3 Policy Attach Points

This section describes each of the OSPFv3 policy attach points and provides a summary of the OSPFv3 attributes and operators.

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

This table summarizes the OSPFv3 attributes and operators per attach points.

Table 10: OSPFv3 Attributes and Operators

Attach Point	Attribute	Match	Set
default-information originate	ospf-metric	—	set
originate	metric-type		set
	tag		set
	rib-has-route	in	
redistribute	destination	in	—
	ospf-metric	—	set
	metric-type		set
	route-type	is	—
	tag	is, eq, ge, le	

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.

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
```

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
  .
```

Nondestructive Editing of Routing Policy

The Nondestructive Editing of Routing Policy changes the default exit behavior under routing policy configuration mode to terminate the configuration.

The default **exit** command acts as end-policy, end-set, or end-if. If the **exit** command is executed under route policy configuration mode, the changes are applied and configuration is updated. This destructs the existing policy. The **rpl set-exit-as-abort** command allows to overwrite the default behavior of the **exit** command under the route policy configuration mode.

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 reentering 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.



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 the software for editing RPL policies:

- Nano (default)
- Emacs
- Vim

Editing Routing Policy Configuration Elements Using Emacs Editor

To edit the contents of a routing policy using the Emacs editor, use the following CLI command in XR 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
```

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
1
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 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 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 mode.

Editing Routing Policy Configuration Elements Using Nano Editor

To edit the contents of a routing policy using the Nano editor, use the following CLI command in XR 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 the software.

Editing Routing Policy Language set elements Using XML

RPL supports editing set elements using XML. Entries can be appended, prepended, or deleted to an existing set without replacing it through XML.

Hierarchical Policy Conditions

The Hierarchical Policy Conditions feature enables the ability to specify a route policy within the "if" statement of another route policy. This ability enables route-policies to be applied for configurations that are based on hierarchical policies.

With the Hierarchical Policy Conditions feature, the software supports Apply Condition policies that can be used with various types of Boolean operators along with various other matching statements.

Apply Condition Policies

Apply Condition policies allow usage of a route-policy within an "if" statement of another route-policy.

Consider route-policy configurations Parent, Child A, and Child B:

```
route-policy ChildA
if destination in (10.10.0.0/16) then
 set local-pref 111
 endif
end-policy
1
route-policy ChildB
if as-path originates-from '222' then
 set community (333:222) additive
 endif
end-policy
!
route-policy Parent
if apply ChildA and apply ChildB then
 set community (333:333) additive
 else
 set community (333:444) additive
endif
end-policy
1
```

In the above scenarios, whenever the policy *Parent* is executed, the decision of the "if" condition in that is selected based on the result of policies *Child A* and *Child B*. The policy *Parent* is equivalent to policy *merged* as given below:

```
route-policy merged
if destination in (10.10.0.0/16) and as-path originates-from '222' then
set local-pref 111
set community (333:222, 333:333) additive
elseif destination in (10.10.0.0/16) then /*Only Policy ChildA is pass */
set local-pref 111
set community (333:444) additive /*From else block */
elseif as-path originates-from '222' then /*Only Policy ChildB is pass */
set community (333:222, 333:444) additive /*From else block */
else
set community (333:444) additive /*From else block */
else
set community (333:444) additive /*From else block */
endif
end-policy
```

Apply Conditions can be used with parameters and are supported on all attach points and on all clients. Hierarchical Apply Conditions can be used without any constraints on a cascaded level.

Existing route policy semantics can be expanded to include this Apply Condition:

```
Route-policy policy_name

If apply policyA and apply policyB then

Set med 100

Else if not apply policyD then

Set med 200

Else

Set med 300

Endif

End-policy
```

Behavior of pass/drop/done RPL Statements for Simple Hierarchical Policies

This table describes the behavior of **pass/drop/done** RPL statements, with a possible sequence for executing the **done** statement for Simple Hierarchical Policies.

Route-policies with simple hierarchical policies	Possible done statement execution sequence	Behavior
pass	pass Continue_list	Marks the prefix as "acceptable" and continues with execution of continue_list statements.
drop	Stmts_list drop	Rejects the route immediately on hitting the drop statement and stops policy execution.
done	Stmts_list done	Accepts the route immediately on hitting the done statement and stops policy execution.
pass followed by done	pass Statement_list done	Exits immediately at the done statement with "accept route".
drop followed by done	drop Statement list done	This is an invalid scenario at execution point of time. Policy terminates execution at the drop statement itself, without going through the statement list or the done statement; the prefix will be rejected or dropped.

Behavior of pass/drop/done RPL Statements for Hierarchical Policy Conditions

This section describes the behavior of **pass/drop/done** RPL statements, with a possible sequence for executing the **done** statement for Hierarchical Policy Conditions.

Terminology for policy execution: "true-path", "false-path", and "continue-path".

```
Route-policy parent

If apply hierarchical_policy_condition then

TRUE-PATH : if hierarchical_policy_condition returns TRUE then this path will

be executed.

Else

FALSE-PATH : if hierarchical_policy_condition returns FALSE then this path will

be executed.

End-if

CONTINUE-PATH : Irrespective of the TRUE/FALSE this path will be executed.

End-policy
```

Hierarchical policy conditions	Possible done statement execution sequence	Behavior
pass	pass Continue_list	Marks the return value as "true" and continues execution within the same policy condition. If there is no statement after " pass ", returns "true".
pass followed by done	pass or set action statement Stmt_list done	Marks the return value as "true" and continues execution till the done statement. Returns "true" to the apply policy condition to take "true-path".
done	Stmt_list without pass or set operation DONE	Returns " false". Condition takes "false-path".
drop	Stmt_list drop Stmt_list	The prefix is dropped or rejected.

Nested Wildcard Apply Policy

The hierarchical constructs of Routing Policy Language (RPL) allows one policy to refer to another policy. The referred or called policy is known as a child policy. The policy from which another policy is referred is called calling or parent policy. A calling or parent policy can nest multiple child policies for attachment to a common set of BGP neighbors. The nested wildcard apply policy allows wildcard (*) based apply nesting. The wildcard operation permits declaration of a generic apply statement that calls all policies that contain a specific defined set of alphanumeric characters, defined on the router.

A wildcard is specified by placing an asterisk (*) at the end of the policy name in an apply statement. Passing parameters to wildcard policy is not supported. The wildcard indicates that any value for that portion of the apply policy matches.

To illustrate nested wildcard apply policy, consider this policy hierarchy:

```
route-policy Nested_Wilcard
apply service_policy_customer*
end-policy
route-policy service_policy_customer_a
if destination in prfx_set_customer_a then
set extcommunity rt (1:1) additive
endif
end-policy
route-policy service_policy_customer_b
if destination in prfx_set_customer_b then
set extcommunity rt (1:1) additive
endif
end-policy
```

```
route-policy service_policy_customer_c
if destination in prfx_set_customer_c then
set extcommunity rt (1:1) additive
endif
end-policy
```

Here, a single parent apply statement (apply service_policy_customer*) calls (inherits) all child polices that contain the identified character string "service_policy_customer". As each child policy is defined globally, the parent dynamically nests the child policies based on the policy name. The parent is configured once and inherits each child policy on demand. There is no direct association between the parent and the child policies beyond the wildcard match statement.

Match Aggregated Route

The Match Aggregated Route feature helps to match BGP aggregated route from the non-aggregated route. BGP can aggregate a group of routes into a single prefix before sending updates to a neighbor. With Match Aggregated Route feature, route policy separates this aggregated route from other routes.

Remove Private AS in Inbound Policy

BGP appends its own as-path before sending out packets to neighbors. When a packet traverses multiple iBGP neighbors, the as-path structure will have many private autonomous systems (AS) in them. The Remove Private AS in Inbound Policy will give the capability to delete those private autonomous systems using RPL route-policy. The **remove as-path private-as** command removes autonomous systems (AS) with AS number 64512 through 65535.



Implementing Static Routes

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 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.

References for Static Routes, on page 195 provides additional conceptual information on static routes.

This module describes how to implement static routes.

- Restrictions for Implementing Static Routes, on page 189
- Configure Static Route, on page 189
- Floating Static Routes, on page 191
- Change Maximum Number of Allowable Static Routes, on page 193
- Default VRF, on page 194
- References for Static Routes, on page 195
- IPv4 Multicast Static Routes, on page 197

Restrictions for Implementing Static Routes

These restrictions apply while implementing Static Routes:

- Static routing to an indirect next hop, (any prefix learnt through the RIB and may be more specific over the AIB), that is part of a local subnet requires configuring static routes in the global table indicating the egress interfaces as next hop. To avoid forward drop, configure static routes in the global table indicating the next-hop IP address to be the next hop.
- Generally, a route is learnt from the AIB in the global table and is installed in the FIB. However, this behavior will not be replicated to leaked prefixes. This could lead to inconsistencies in forwarding behavior.

Configure Static Route

Static routes are entirely user configurable and can point to a next-hop interface, next-hop IP address, or both. In the 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.

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. Use the commit or end command.

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router static
	Example:
	RP/0/RP0/CPU0:router(config)# router static
	Enters static route configuration mode.
Step 3	vrf vrf-name
	Example:
	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	address-family { ipv4 ipv6 } { unicast multicast }
	Example:
	RP/0/RP0/CPU0:router(config-static-vrf)# address-family ipv4 unicast
	Enters address family mode.
Step 5	prefix mask [vrf vrf-name] { ip-address interface-type interface-instance } [distance] [description text] [tag tag] [permanent]
	Example:
	RP/0/RP0/CPU0:router(config-static-vrf-afi)# 10.0.0.0/8 172.20.16.6 110

Configures an administrative distance of 110.

• 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.

Step 6 Use the **commit** or **end** command.

commit—Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No -Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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
```

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.



Note By default, static routes have smaller administrative distances than dynamic routes, so static routes are preferred to dynamic routes.

Configure 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. Use the commit or end command.

DETAILED STEPS

Step 1	configure
	Example:
	RP/0/RP0/CPU0:router# configure
	Enters mode.
Step 2	router static
	Example:
	RP/0/RP0/CPU0:router(config) # router static
	Enters static route configuration mode.
Step 3	vrf vrf-name
	Example:
	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	address-family { ipv4 ipv6 } { unicast multicast }
	Example:
	RP/0/RP0/CPU0:router(config-static-vrf)# address-family ipv6 unicast
	Enters address family mode.
Step 5	prefix mask [vrf vrf-name] { ip-address interface-type interface-instance } [distance] [description text] [tag tag] [permanent]
	Example:
	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	Use the commit or end command.
	commit —Saves the configuration changes and remains within the configuration session.
	end —Prompts user to take one of these actions:
	• Yes — Saves configuration changes and exits the configuration session.
	• No —Exits the configuration session without committing the configuration changes.

• Cancel — Remains in the configuration session, without committing the configuration changes.

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
```

Change Maximum Number of Allowable Static Routes

This task explains how to change the maximum number of allowable static routes.

Before you begin



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. Use the commit or end command.

DETAILED STEPS

Step 1 configure Example: RP/0/RP0/CPU0:router# configure Enters mode.

Step 2 router static

Example:

RP/0/RP0/CPU0:router(config) # router static
Enters static route configuration mode.

Step 3 maximum path { ipv4 | ipv6 } value Example:

RP/0/RP0/CPU0:router(config-static) # maximum path ipv4 10000

Changes the maximum number of allowable static routes.

- Specify IPv4 or IPv6 address prefixes.
- Specify the maximum number of static routes for the given address family. The range is from 1 to 140000.
- This example sets the maximum number of static IPv4 routes to 10000.

Step 4 Use the **commit** or **end** command.

commit ---Saves the configuration changes and remains within the configuration session.

end —Prompts user to take one of these actions:

- Yes Saves configuration changes and exits the configuration session.
- No —Exits the configuration session without committing the configuration changes.
- Cancel Remains in the configuration session, without committing the configuration changes.

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
```

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.



Note 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.

References for Static Routes

The following topics provide additional conceptual information on static routes:

- Static Route Functional Overview, on page 195
- Default Administrative Distance, on page 195
- Directly Connected Routes, on page 196
- Floating Static Routes , on page 191
- Fully Specified Static Routes , on page 196
- Recursive Static Routes, on page 196

Static Route Functional Overview

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.

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 TenGigE 0/0/0/0:

```
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 TenGigE 0/0/0/0
```

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.

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.



Note By default, static routes have smaller administrative distances than dynamic routes, so static routes are preferred to dynamic routes.

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 TenGigE 0/0/0/0 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.

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 is re-inserted in the routing table.

Dynamic ECMP

The dynamic ECMP (equal-cost multi-path) for IGP (Interior Gateway Protocol) prefixes feature supports dynamic selection of ECMP paths ranging from 1 to 64 IGP paths. ECMP for non-recursive prefixes is dynamic. This feature enables loadbalancing support in hardware among egress links.

The dynamic ECMP (equal-cost multi-path) for IGP (Interior Gateway Protocol) prefixes feature supports dynamic selection of ECMP paths ranging from 1 to 64 IGP paths. ECMP for non-recursive prefixes is dynamic.

This feature enables loadbalancing support in hardware among egress links.


IPv4 Multicast Static Routes

IP multicast static routes (mroutes) allow you to have multicast paths diverge from the unicast paths. When using Protocol Independent Multicast (PIM), the router expects to receive packets on the same interface where it sends unicast packets back to the source. This expectation is beneficial if your multicast and unicast topologies are congruent. However, you might want unicast packets to take one path and multicast packets to take another.

The most common reason for using separate unicast and multicast paths is tunneling. When a path between a source and a destination does not support multicast routing, configuring two routers with a GRE tunnel

between them is the solution. In the figure below, each unicast router (UR) supports unicast packets only; each multicast router (MR) supports multicast packets.

Figure 5: Tunnel for Multicast Packets



In the figure, the source delivers multicast packets to destination by using MR 1 and MR 2. MR 2 accepts the multicast packet only if it predicts it can reach source over the tunnel. If this situation is true, when the destination sends unicast packets to the source, MR 2 sends them over the tunnel. The check that MR2 can reach the source over the tunnel is a Reverse Path Forwarding (RPF) check, and the static mroute allows the check to be successful when the interface, on which the multicast packet arrives, is not the unicast path back to the source. Sending the packet over the tunnel could be slower than natively sending it through UR 2, UR 1, and MR 1.

A multicast static route allows you to use the configuration in the above figure by configuring a static multicast source. The system uses the configuration information instead of the unicast routing table to route the traffic. Therefore, multicast packets can use the tunnel without having the unicast packets use the tunnel. Static mroutes are local to the router they are configured on and not advertised or redistributed in any way to any other router.

Configure Multicast Static Routes

The following example shows how to configure multiple static routes in IPv4 and IPv6 address family configuration modes:

```
/* Enables a static routing process */
Router (config) # router static
/* Configures the IPv4 address-family for the unicast topology with a destination prefix.
*/
Router(config-static) # address-family ipv4 unicast
Router(config-static-afi) # 10.1.1.0/24 198.51.100.1
Router (config-static-afi) # 223.255.254.254/32 203.0.113.1
Router(config-static-afi) # exit
/* Configures the IPv4 address-family for the multicast topology with a destination prefix.
 */
Router (config-static) # address-family ipv4 multicast
Router(config-static-afi)# 198.51.100.20/32 209.165.201.0
Router(config-static-afi) # 192.0.2.10/32 209.165.201.0
Router(config-static-afi)# exit
/* Enable the address family IPv4 and IPv6 multicast on the next hop interface. */
Router(config) # interface TenGigE 0/0/0/12
Router(config-if)# address-family ipv4 multicast
Router(config-if)# address-family ipv6 multicast
```

L

Running Configuration

```
router static
address-family ipv4 unicast
10.1.1.0/24 198.51.100.1
223.255.254.254/32 203.0.113.1
!
address-family ipv4 multicast
198.51.100.20/32 209.165.201.0
192.0.2.10/32 209.165.201.0
!
interface TenGigE 0/0/0/12
address-family ipv4 multicast
address-family ipv6 multicast
```

Verification

Verify the IPv4 multicast routes.

```
show route ipv4 multicast
```

```
Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
     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, G - DAGR, l - LISP
      A - access/subscriber, a - Application route
      M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup path
Gateway of last resort is 10.1.1.20 to network 0.0.0.0
i*L1 0.0.0.0/0 [115/10] via 10.1.1.20, 00:41:12, TenGigE0/0/0/6
     10.1.1.0/24 is directly connected, 00:41:12, TenGigE0/0/0/0
С
     10.1.1.10/32 is directly connected, 00:41:12, TenGigE0/0/0/0
T.
S
     172.16.2.10/32 [1/0] via 198.51.100.20, 00:41:12
i L1 172.16.3.1/32 [115/20] via 198.51.100.20, 00:41:12, TenGigE0/0/0/12
i L1 192.0.2.1/24 [115/20] via 198.51.100.20, 00:41:12, TenGigE0/0/0/1
```



Implementing Integrated Routing and Bridging

IRB provides the ability to route between a bridge group and a routed interface using a BVI. The BVI is a virtual interface within the router that acts like a normal routed interface. A BVI is associated with a single bridge domain and represents the link between the bridging and the routing domains on the router. To support receipt of packets from a bridged interface that are destined to a routed interface, the BVI must be configured with the appropriate IP addresses and relevant Layer 3 attributes.

- Bridge-Group Virtual Interface, on page 201
- Sub-Interfaces with BVI interface in Bridge Domain, on page 202
- Supported Features on a BVI, on page 202
- BVI Interface and Line Protocol States, on page 202
- Prerequisites for Configuring IRB, on page 203
- Restrictions for Configuring IRB, on page 203
- How to Configure IRB, on page 203
- Additional Information on IRB, on page 212
- Packet Flows Using IRB, on page 212
- Configuration Examples for IRB, on page 214

Bridge-Group Virtual Interface

The BVI is a virtual interface within the router that acts like a normal routed interface. The BVI does not support bridging itself, but acts as a gateway for the corresponding bridge-domain to a routed interface within the router.

Aside from supporting a configurable MAC address, a BVI supports only Layer 3 attributes, and has the following characteristics:

- Uses a MAC address taken from the local chassis MAC address pool, unless overridden at the BVI interface.
- Is configured as an interface type using the **interface bvi** command and uses an IPv4 address that is in the same subnet as the hosts on the segments of the bridged domain. The BVI also supports secondary addresses.
- The BVI identifier is independent of the bridge-domain identifier. These identifiers do not need to correlate like they do in Cisco IOS software.
- Is associated to a bridge group using the **routed interface bvi** command.

Sub-Interfaces with BVI interface in Bridge Domain

Support for Layer 2 (L2) sub-interfaces to work with BVI interface in a bridge domain has been added in the Cisco IOS XR Software Release 6.3.1. The L2 sub-interfaces must have a rewrite configuration to pop the outer tag in Ingress direction and push the outer tag on the packet in egress direction. The outer tag that is to be pushed in egress direction is based on the interface encapsulation. The support is added only for single tag sub-interfaces. Double tagged sub-interfaces that work with BVI interface in a bridge domain are not supported.

Supported Features on a BVI

- These interface commands are supported on a BVI:
 - arp purge-delay
 - arp timeout
 - bandwidth (The default is 10 Gbps and is used as the cost metric for routing protocols for the BVI)
 - ipv4
 - ipv6
 - mac-address
 - mtu (The default is 1500 bytes)
 - shutdown
- The BVI supports IP helper addressing and secondary IP addressing.

BVI Interface and Line Protocol States

Like typical interface states on the router, a BVI has both an Interface and Line Protocol state.

- The BVI interface state is Up when the following occurs:
 - The BVI interface is created.
 - The bridge-domain that is configured with the **routed interface bvi** command has at least one available active bridge port .
- These characteristics determine when the the BVI line protocol state is up:
 - The bridge-domain is in Up state.
 - The BVI IP address is not in conflict with any other IP address on another active interface in the router.

Prerequisites for Configuring IRB

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Before configuring IRB, be sure that these tasks and conditions are met:

- Know the IP addressing and other Layer 3 information to be configured on the bridge virtual interface (BVI).
- Complete MAC address planning if you decide to override the common global MAC address for all BVIs.
- Be sure that the BVI network address is being advertised by running static or dynamic routing on the BVI interface.

Restrictions for Configuring IRB

Before configuring IRB, consider these restrictions:

- Only one BVI can be configured in any bridge domain.
- The same BVI can not be configured in multiple bridge domains.
- The following areas are *not* supported on the Layer2 bridging (with BVI):
 - Access Control Lists (ACLs). However, Layer 2 ACLs can be configured on each Layer 2 port of the bridge domain.
 - Static mac entry configuration in Bridge.
 - Mac ageing configuration at global config mode.
 - MAC Learning Disable.
 - Port-channel sub-interface as bridge member.
 - Physical sub-interface as bridge member.
 - Vlan rewrite.
- QOS configuration on BVI interface is not supported.
- VRF on BVI interface is not supported.

How to Configure IRB

This section includes the following configuration tasks:

Configuring the Bridge Group Virtual Interface

To configure a BVI, complete the following steps.

Configuration Guidelines

Consider the following guidelines when configuring the BVI:

- The BVI must be assigned an IPv4 or IPv6 address that is in the same subnet as the hosts in the bridged segments.
- If the bridged network has multiple IP networks, then the BVI must be assigned secondary IP addresses for each network.

SUMMARY STEPS

- 1. configure
- 2. interface bvi identifier
- **3.** ipv4 address *ipv4-address mask* [secondary] ipv6 address *ipv6-prefix/prefix-length* [eui-64] [route-tag *route-tag value*]
- 4. arp purge-delay seconds
- 5. arp timeout seconds
- 6. bandwidth *rate*
- 7. mac-address value1.value2.value3
- 8. mtu bytes
- 9. end or commit

DETAILED STEPS

Step 1 configure Example: RP/0/RP0/CPU0:router# configure Enters the global configuration mode. Step 2 interface bvi identifier Example: RP/0/RP0/CPU0:router(config)# interface bvi 1

MI/0/MI0/0100.100001 (coming) # incertace DVI 1

Specifies or creates a BVI, where *identifier* is a number from 1 to 65535.

 Step 3
 ipv4 address ipv4-address mask [secondary] ipv6 address ipv6-prefix/prefix-length [eui-64] [route-tag route-tag value]

 Example:

RP/0/RP0/CPU0:router(config-if) # ipv4 address 10.10.0.4 255.255.255.0

Specifies a primary or secondary IPv4 address or an IPv6 address for an interface.

Step 4 arp purge-delay seconds

Example:

RP/0/RP0/CPU0:router(config-if)#arp purge-delay 120

(Optional) Specifies the amount of time (in *seconds*) to delay purging of Address Resolution Protocol (ARP) table entries when the interface goes down.

The range is 1 to 65535. By default purge delay is not configured.

Step 5 arp timeout seconds

Example:

RP/0/RP0/CPU0:router(config-if)# arp timeout 12200

(Optional) Specifies how long dynamic entries learned on the interface remain in the ARP cache.

The range is 30 to 2144448000 seconds. The default is 14,400 seconds (4 hours).

Step 6 bandwidth *rate*

Example:

RP/0/RP0/CPU0:router(config-if)# bandwidth 1000000

(Optional) Specifies the amount of bandwidth (in kilobits per second) to be allocated on the interface. This number is used as the cost metric in routing protocols for the BVI.

The range is 0 to 4294967295. The default is 10000000 (10 Gbps).

Step 7 mac-address value1.value2.value3

Example:

RP/0/RP0/CPU0:router(config-if)# mac-address 1111.2222.3333

(Optional) Specifies the 48-bit MAC address for the BVI as three dotted-hexadecimal values, and overrides use of the default MAC address. The range for each value is 0000 to ffff. A MAC address of all 0s is not supported.

Step 8 mtu bytes

Example:

RP/0/RP0/CPU0:router(config-if) # mtu 2000

(Optional) Specifies the maximum transmission unit (MTU) size for packets on the interface. The range is 64 to 65535. The default is 1514.

Step 9 end or commit

Example:

RP/0/RP0/CPU0:router(config-if) # end

or

RP/0/RP0/CPU0:router(config-if) # commit

Saves configuration changes.

• When you issue the **end** command, the system prompts you to commit changes:

Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:

Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

 Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring the Layer 2 AC Interfaces

To configure the Layer 2 AC interfaces for routing by a BVI, complete the following steps.

SUMMARY STEPS

- 1. configure
- 2. interface [HundredGigE | TenGigE] interface-path-idl2transport
- 3. end or commit

DETAILED STEPS

Step 1 configure

Example:

RP/0/RP0/CPU0:router# configure

Enters global configuration mode.

Step 2 interface [HundredGigE | TenGigE] interface-path-idl2transport

Example:

RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/3 l2transport

Enables Layer 2 transport mode on interface, where *interface-path-id* is specified as the *rack/slot/module/port* location of the interface.

Step 3 end or commit

Example:

RP/0/RP0/CPU0:router(config-if) # end

or

RP/0/RP0/CPU0:router(config-if) # commit

Saves configuration changes.

• When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring a Bridge Group and Assigning Interfaces to a Bridge Domain

To configure a bridge group and assign interfaces to a bridge domain, complete the following steps.

SUMMARY STEPS

- 1. configure
- 2. l2vpn
- **3.** bridge group bridge-group-name
- 4. bridge-domain bridge-domain-name
- 5. interface [HundredGigE | TenGigE] interface-path-id
- 6. end or commit

DETAILED STEPS

 Step 1
 configure

 Example:
 RP/0/RP0/CPU0:router# configure

 Enters global configuration mode.

 Step 2
 I2vpn

 Example:

 RP/0/RP0/CPU0:router(config)# 12vpn

Enters L2VPN configuration mode.

Step 3 bridge group bridge-group-name

Example:

RP/0/RP0/CPU0:router(config-l2vpn)# bridge group 10

Creates a bridge group and enters L2VPN bridge group configuration mode.

Step 4 bridge-domain bridge-domain-name Example:

RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain BD 1

Creates a bridge domain and enters L2VPN bridge group bridge domain configuration mode.

Step 5 interface [HundredGigE | TenGigE] interface-path-id

Example:

RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# interface HundredGigE 0/1/0/3

Associates the 100-Gigabit Ethernet or 10-Gigabit Ethernet interface with the specified bridge domain and enters L2VPN bridge group bridge domain attachment circuit configuration mode, where *interface-path-id* is specified as the *rack/slot/module/port* location of the interface.

Repeat this step for as many interfaces as you want to associate with the bridge domain.

Step 6 end or commit

Example:

RP/0/RP0/CPU0:router(config=l2vpn-bg-bd-ac)# end

or

RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac) # commit

Saves configuration changes.

• When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

• Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Configure the Layer 2 AC Sub-Interfaces

To configure the Layer 2 AC sub-interfaces for routing by a BVI, complete the following steps.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface TenGigE0/0/0/0.1 l2transport
RP/0/RSP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RSP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
```

Running Configuration

```
configure
interface TenGigE0/0/0/0.1 l2transport
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric
```

Verification

Verify configuration of layer 2 AC sub-interface.

RP/0/RP0/CPU0:ios# show running interface TenGigE 0/0/0/0.1
Wed Jul 26 13:40:50.173 UTC
interface TenGigE0/0/0/0.1 l2transport

```
encapsulation dotlq 100
rewrite ingress tag pop 1 symmetric
```

Configure a Bridge Group and Assign Sub-Interfaces to a Bridge Domain

To configure a bridge group and assign a sub-interface to a bridge domain, complete the following steps.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# 12vpn
RP/0/RSP0/CPU0:router(config-12vpn)# bridge group 1
RP/0/RSP0/CPU0:router(config-12vpn-bg)# bridge-domain bd1
RP/0/RSP0/CPU0:router(config-12vpn-bg-bd)# bridge-domain bd1
RP/0/RSP0/CPU0:router(config-12vpn-bg-bd-ac)# interface TenGigE 0/0/0/0.1
RP/0/RSP0/CPU0:router(config-12vpn-bg-bd-ac)# exit
RP/0/RSP0/CPU0:router(config-12vpn-bg-bd-ac)# exit
```

Running Configuration

configure 12vpn

```
bridge group 1
bridge-domain BD1
interface TenGigE0/0/0/0.1
!
routed interface BVI1
!
```

Verification

Verify the configuration of the bridge group and assignment of a sub-interface to bridge-domain.

```
RP/0/RP0/CPU0:ios#show running l2vpn
Wed Jul 26 13:38:32.830 UTC
l2vpn
bridge group 1
    bridge-domain BD1
    interface TenGigE 0/0/0/0.1
    !
    routed interface BVI1
    !
```

Associating the BVI as the Routed Interface on a Bridge Domain

To associate the BVI as the routed interface on a bridge domain, complete the following steps.

SUMMARY STEPS

- 1. configure
- 2. l2vpn
- 3. bridge group bridge-group-name
- 4. bridge-domain bridge-domain-name
- 5. routed interface bvi identifier
- 6. end or commit

DETAILED STEPS

Step 1	configure Example:
	RP/0/RP0/CPU0:router# configure Enters global configuration mode.
Step 2	l2vpn Example:
	RP/0/RP0/CPU0:router(config)# 12vpn

Enters L2VPN configuration mode.

Step 3	bridge group bridge-group-name
	Example:
	RP/0/RP0/CPU0:router(config-l2vpn)# bridge group BG_test
	Creates a bridge group and enters L2VPN bridge group configuration mode.
Step 4	bridge-domain bridge-domain-name
	Example:
	RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain 1
	Creates a bridge domain and enters L2VPN bridge group bridge domain configuration mode.
Step 5	routed interface bvi identifier
	Example:
	RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1
	Associates the specified BVI as the routed interface for the interfaces assigned to the bridge domain.
Step 6	end or commit
	Example:
	RP/0/RP0/CPU0:router(config-12vpn-bg-bd)# end
	or
	RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
	Saves configuration changes.
	• When you issue the end command, the system prompts you to commit changes:

Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Displaying Information About a BVI

To display information about BVI status and packet counters, use the following commands:


show interfaces byi <i>identifier</i> [accounting brief description detail]	Displays interface status, line protocol state, and packet counters for the specified BVI.
show adjacency bvi identifier [detail remote]	Displays packet and byte transmit counters per adjacency to the specified BVI.
show l2vpn bridge-domain detail	Displays the reason that a BVI is down.

Additional Information on IRB

Packet Flows Using IRB

This figure shows a simplified functional diagram of an IRB implementation to describe different packet flows between Host A, B, and C. In this example, Host C is on a network with a connection to the same router. In reality, another router could be between Host C and the router shown.

Figure 6: IRB Packet Flows Between Hosts



When IRB is configured on a router, the following processing happens:

- ARP requests are resolved between the hosts and BVI that are part of the bridge domain.
- All packets from a host on a bridged interface go to the BVI if the destination MAC address matches the BVI MAC address. Otherwise, the packets are bridged.
- For packets destined for a host on a routed network, the BVI forwards the packets to the routing engine before sending them out a routed interface.
- All packets either from or destined to a host on a bridged interface go to the BVI first (unless the packet is destined for a host on the bridge domain).
- For packets that are destined for a host on a segment in the bridge domain that come in to the router on a routed interface, the BVI forwards the packet to the bridging engine, which forwards it through the appropriate bridged interface.

Packet Flows When Host A Sends to Host B on the Bridge Domain

When Host A sends data to Host B in the bridge domain on the 10.10.0.0 network, no routing occurs. The hosts are on the same subnet and the packets are bridged between their segment interfaces on the router.

Packet Flows When Host A Sends to Host C From the Bridge Domain to a Routed Interface

Using host information from this figure, the following occurs when Host A sends data to Host C from the IRB bridging domain to the routing domain:

- Host A sends the packet to the BVI (as long any ARP request the is resolved between the host and the BVI). The packet has the following information:
 - Source MAC address of host A.
 - Destination MAC address of the BVI.
- Since Host C is on another network and needs to be routed, the BVI forwards the packet to the routed interface with the following information:
 - IP source MAC address of Host A (10.10.0.2) is changed to the MAC address of the BVI (10.10.0.4).
 - IP destination address is the IP address of Host C (10.20.0.3).
- Interface 10.20.0.2 sees receipt of a packet from the routed BVI 10.10.0.4. The packet is then routed through interface 10.20.0.2 to Host C.

Packet Flows When Host C Sends to Host B From a Routed Interface to the Bridge Domain

Using host information from this figure, the following occurs when Host C sends data to Host B from the IRB routing domain to the bridging domain:

- The packet comes into the routing domain with the following information:
 - MAC source address—MAC of Host C.
 - MAC destination address—MAC of the 10.20.0.2 ingress interface.
 - IP source address—IP address of Host C (10.20.0.3).
 - IP destination address—IP address of Host B (10.10.0.3).
- When interface 10.20.0.2 receives the packet, it looks in the routing table and determines that the packet needs to be forwarded to the BVI at 10.10.0.4.
- The routing engine captures the packet that is destined for the BVI and forwards it to the BVI's corresponding bridge domain. The packet is then bridged through the appropriate interface if the destination MAC address for Host B appears in the bridging table, or is flooded on all interfaces in the bridge group if the address is not in the bridging table.

Configuration Examples for IRB

This section provides the following configuration examples:

Basic IRB Configuration: Example

The following example shows how to perform the most basic IRB configuration:

```
! Configure the BVI and its IPv4 address
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#interface bvi 1
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.10.0.4 255.255.255.0
RP/0/RP0/CPU0:router(config-if))# exit
! Configure the Layer 2 AC interface
RP/0/RP0/CPU0:router(config) #interface HundredGigE 0/1/0/0 12transport
RP/0/RP0/CPU0:router(config-if))# exit
! Configure the L2VPN bridge group and bridge domain and assign interfaces
T
RP/0/RP0/CPU0:router(config) #12vpn
RP/0/RP0/CPU0:router(config-l2vpn) #bridge group 10
RP/0/RP0/CPU0:router(config-l2vpn-bg)#bridge-domain 1
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-if) # exit
! Associate a BVI to the bridge domain
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

IPv4 Addressing on a BVI Supporting Multiple IP Networks: Example

The following example shows how to configure secondary IPv4 addresses on a BVI that supports bridge domains for the 10.10.10.0/24, 10.20.20.0/24, and 10.30.30.0/24 networks. In this example, the BVI must have an address on each of the bridge domain networks:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#interface bvi 1
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.10.10.4 255.255.255.0
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.20.20.4 255.255.255.0 secondary
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.30.30.4 255.255.255.0 secondary
RP/0/RP0/CPU0:router(config-if)# commit
```



Implementing BFD

Bidirectional forwarding detection (BFD) provides low-overhead, short-duration detection of failures in the path between adjacent forwarding engines. BFD allows a single mechanism to be used for failure detection over any media and at any protocol layer, with a wide range of detection times and overhead. The fast detection of failures provides immediate reaction to failure in the event of a failed link or neighbor.



Note Cisco NCS 5000 Routers do not support BFD dampening. If it is enabled by default, disable it first before configuring any session. If dampening is disabled after configuring the BFD session, reload the device for it to take effect.

Tip You can programmatically configure BFD and retrieve operational data using <code>openconfig-bfd.yang</code> OpenConfig data model. To get started with using data models, see the *Programmability Configuration Guide*.

- BFD over Bundle, on page 215
- BFD over Bundle and BFD over Logical Bundle, on page 218
- Enabling BFD Sessions on Bundle Members, on page 218
- Specifying the BFD Destination Address on a Bundle, on page 219
- Configuring the Minimum Thresholds for Maintaining an Active Bundle, on page 220
- Configuring BFD Packet Transmission Intervals and Failure Detection Times on a Bundle, on page 222
- Configure BFD over Bundles IETF Mode Support on a Per Bundle Basis, on page 223

BFD over Bundle

BFD over Bundle feature enables BFD sessions to monitor the status of individual bundle member links. BFD notifies the bundle manager immediately when one of the member links goes down, and reduces the bandwidth used by the bundle.

Restrictions

The following are the restrictions in using BFD over Bundle feature:

- It is only supported in IETF mode.
- It is only supported on main bundle interface; it is not supported on bundle sub-interfaces.

- It is not supported on routing protocols, such as OSPF, ISIS, and BGP.
- When BFD timer is configured to 3.3 ms, which is the most aggressive timer, 256 sessions can be brought up.
- If BFD timer is configured to greater than 100 ms, 300 BFD sessions can be brought up simultaneously.
- BFD echo mode and encryption is not supported.
- BFD dampening is not supported.

Configure BFD over Bundle

Configuring BFD over bundle involves the following steps:

- Enable and Disable IPv6 checksum calculations for BFD on a router
- Specify the mode, BFD packet transmission intervals, and failure detection times on a bundle

Note

Repeat the same configuration steps in the destination router.

```
/* Enable and Disable IPv6 checksum calculations for BFD on a router. ^{\prime}
Router(config-if) # bfd
Router(config-bfd-if) # ipv6 checksum disable
Router(config-bfd-if) # dampening disable
Router(config-bfd-if) # commit
/* Specify the mode, BFD packet transmission intervals, and failure detection times on a
bundle */
Router(config) # interface Bundle-Ether 3739
Router(config-if) # bfd mode ietf
Router(config-if) # bfd address-family ipv4 multiplier 3
Router(config-if)# bfd address-family ipv4 destination 10.23.1.2
Router(config-if) # bfd address-family ipv4 fast-detect
Router(config-if)# bfd address-family ipv4 minimum-interval 100
Router(config-if) # bfd address-family ipv6 multiplier 3
Router(config-if)# bfd address-family ipv6 destination 2001:DB8:1::2
Router(config-if) # bfd address-family ipv6 fast-detect
Router(config-if)# bfd address-family ipv6 minimum-interval 100
Router(config-if) # ipv4 address 10.23.1.1 255.255.255.252
Router(config-if) # ipv6 address 2001:DB8:1::2/120
Router(config-if) # load-interval 30
Router(config-if) # commit
Router(config) # interface TenGigE 0/0/0/0
Router(config-if) # bundle id 3739 mode active
```

Running Configuration

```
bfd
ipv6 checksum disable
dampening disable!
!
```

```
interface Bundle-Ether3739
bfd mode ietf
bfd address-family ipv4 multiplier 3
bfd address-family ipv4 destination 10.23.1.2
bfd address-family ipv4 fast-detect
bfd address-family ipv4 minimum-interval 100
bfd address-family ipv6 multiplier 3
bfd address-family ipv6 destination 2001:DB8:1::2
bfd address-family ipv6 fast-detect
bfd address-family ipv6 minimum-interval 100
ipv4 address 10.23.1.1 255.255.255.252
ipv6 address 2001:DB8:1::2/120
load-interval 30
1
interface TenGigE 0/0/0/0
bundle id 3739 mode active
```

Verification

The following show command outputs displays the status of BFD sessions on bundle members:

/* Verify the details of the IPv4 BFD session in the source router. */

Router# show bfd session

 Interface
 Dest Addr
 Local det time(int*mult)
 State
 Echo
 Async
 H/W
 NPU

 ---- ---- ---- ---- ---- --- --- ---

 Te0/0/0/0
 10.23.1.2
 0s(0s*0)
 300ms(100ms*3)
 UP
 Yes
 0/RP0/CPU0

 BE3739
 10.23.1.2
 n/a
 n/a
 UP
 No
 n/a

/* Verify the details of the IPv4 BFD session in the destination router. $^{\prime\prime}$

Router# show bfd session

Interface	Dest Addr	Local det	time(int*mult)	State	Echo	Async	H/W	NPU
Te0/6/0/0	10.23.1.1	0s(0s*0)	300ms(100ms*3)	UP	No	n/a		
BE3739	10.23.1.1	n/a	n/a	UP	No	n/a		

/* Verify the details of the IPv6 BFD session in the source router. $^{*/}$

Router# show bfd ipv6 session

Interface	Dest Addr	Local det	time(int*mult)	State	H/W	NPU	Echo	Async
Te0/0/0/0	10:23:1::2	Yes		0/RP0/0s	(0s*0)	00ms(100ms*3)	UP	
BE3739	10:23:1::2	No		n/a	n/a	n/a	UP	

/* Verify the details of the IPv6 BFD session in the destination router. */

Router# show bfd ipv6 session

Interface	Dest Addr	Local det	time(int*mult)	State	H/W	NPU	Echo	Async
Te0/6/0/0	10:23:1::1	No	n/a	0s(0s*0)		300ms(100ms*3)	UP	
BE3739	10:23:1::1	No	n/a	n/a		n/a	UP	

BFD over Bundle and BFD over Logical Bundle

Link Aggregation Control Protocol (LACP) allows a network device to negotiate an automatic bundling of links by sending LACP packets to their directly connected peer. LACP provides a keep-alive mechanism for the link members. While the default keep-alive is 30s, it is configurable to up to 1s. LACP can detect failures on a per-physical-member link. However, the LACP timers do not fulfill the criteria of current fast convergence requirements.

Differences between BFD over Bundle and BFD over Logical Bundle

BFD over Bundle (BoB) (RFC 7130) has a BFD session on each bundle member. The client is the bundle manager. If a BFD session goes down on a specific member link, the whole bundle interface goes down. That is, when the member link goes down, the number of available links falls below the required minimum. Hence the routing session is brought down.

BFD over Logical Bundle (BLB) (RFC 5880) treats a bundle interface with all its members as a single interface. BLB is a multipath (MP) single-hop session. If BLB is configured on a bundle there is only one single BFD session that is active. This implies that only one bundle member is being monitored by BFD at any given time. The client is one of the routing protocols. When BFD detects a failure, the client brings down the routing session.

The mode (BoB or BLB) is determined by how you configure BFD:

- You can enable BoB by configuring BFD under a Bundle-Ether interface.
- You can enable BLB by configuring BFD under a Bundle-Ether interface on a routing client.

Link Aggregation Control Protocol (LACP) allows a network device to negotiate an automatic bundling of links by sending LACP packets to their directly connected peer. LACP provides a keep-alive mechanism for the link members. While the default keep-alive is 30s, it is configurable to up to 1s. LACP can detect failures on a per-physical-member link. However, the LACP timers do not fulfill the criteria of current fast convergence requirements.

Enabling BFD Sessions on Bundle Members

To enable BFD sessions on bundle member links, complete these steps:

SUMMARY STEPS

- 1. configure
- 2. interface Bundle-Ether bundle-id
- 3. bfd address-family ipv4 fast-detect
- 4. Use the commit or end command.

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	

	Command or Action	Purpose
	RP/0/RP0/CPU0:router# configure	
Step 2	interface Bundle-Ether bundle-id Example:	Enters interface configuration mode for the specified bundle ID.
	RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1	
Step 3	bfd address-family ipv4 fast-detect	Enables IPv4 BFD sessions on bundle member links.
	Example:	
	<pre>RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 fast-detect</pre>	7
Step 4	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Specifying the BFD Destination Address on a Bundle

To specify the BFD destination address on a bundle, complete these steps:

SUMMARY STEPS

- 1. configure
- **2.** interface Bundle-Ether bundle-id
- 3. bfd address-family ipv4 destination *ip-address*
- 4. Use the commit or end command.

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	

	Command or Action	Purpose	
Step 2	interface Bundle-Ether bundle-id Example:	Enters interface configuration mode for the specified bundle ID.	
	<pre>RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1</pre>		
Step 3	bfd address-family ipv4 destination <i>ip-address</i>	Specifies the primary IPv4 address assigned to the bundle	
	Example:	interface on a connected remote system, where <i>ip-addre</i> is the 32-bit IP address in dotted-decimal format (A.B.C.	
	<pre>RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 destination 10.20.20.1</pre>		
Step 4	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.	
		end —Prompts user to take one of these actions:	
		• Yes — Saves configuration changes and exits the configuration session.	
		 No —Exits the configuration session without committing the configuration changes. 	
		• Cancel —Remains in the configuration session, without committing the configuration changes.	

Configuring the Minimum Thresholds for Maintaining an Active Bundle

The bundle manager uses two configurable minimum thresholds to determine whether a bundle can be brought up or remain up, or is down, based on the state of its member links.

- · Minimum active number of links
- Minimum active bandwidth available

Whenever the state of a member changes, the bundle manager determines whether the number of active members or available bandwidth is less than the minimum. If so, then the bundle is placed, or remains, in DOWN state. Once the number of active links or available bandwidth reaches one of the minimum thresholds, then the bundle returns to the UP state.

To configure minimum bundle thresholds, complete these steps:

SUMMARY STEPS

- 1. configure
- 2. interface Bundle-Ether bundle-id
- 3. bundle minimum-active bandwidth kbps
- 4. bundle minimum-active links links

5. Use the **commit** or **end** command.

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	
Step 2	interface Bundle-Ether <i>bundle-id</i> Example:	Enters interface configuration mode for the specified bundle ID.
	RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1	
Step 3	bundle minimum-active bandwidth kbps	Sets the minimum amount of bandwidth required before a
	Example:	bundle can be brought up or remain up. The range is from 1 through a number that varies depending on the platform
	RP/0/RP0/CPU0:router(config-if)# bundle minimum-active bandwidth 580000	and the bundle type.
Step 4	bundle minimum-active links links	Sets the number of active links required before a bundle
	Example:	can be brought up or remain up. The range is from 1 to 32. Note When BFD is started on a bundle that is
	<pre>RP/0/RP0/CPU0:router(config-if)# bundle minimum-active links 2</pre>	already active, the BFD state of all the existing active members is known.
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Configuring BFD Packet Transmission Intervals and Failure Detection Times on a Bundle

BFD asynchronous packet intervals and failure detection times for BFD sessions on bundle member links are configured using a combination of the **bfd address-family ipv4 minimum-interval** and **bfd address-family ipv4 multiplier** interface configuration commands on a bundle.

The BFD control packet interval is configured directly using the **bfd address-family ipv4 minimum-interval** command. The failure detection times are determined by a combination of the interval and multiplier values in these commands.

To configure the minimum transmission interval and failure detection times for BFD asynchronous mode control packets on bundle member links, complete these steps:

SUMMARY STEPS

- 1. configure
- 2. interface Bundle-Ether bundle-id
- 3. bfd address-family ipv4 minimum-interval milliseconds
- 4. bfd address-family ipv4 multiplier multiplier
- 5. Use the commit or end command.

	Comman	nd or Action	Purpose		
Step 1	configu	re	Enters mode.		
	Example):			
	RP/0/RP	0/CPU0:router# configure			
Step 2	interfac	e Bundle-Ether bundle-id	Enters interface configuration mode for the specified bundle		
	Example):	ID.		
	RP/0/RP Bundle-	0/CPU0:router(config)# interface Ether 1			
Step 3	bfd add	ress-family ipv4 minimum-interval milliseconds			
	Example	2:			
		0/CPU0:router(config-if)#bfd address-family inimum-interval 2000	7		
	Note	Specifies the minimum interval, in milliseconds, for asynchronous mode control packets on IPv4 BFD sessions on bundle member links. The range is from 4 to 30000.			

	Command or Action	Purpose
Step 4	bfd address-family ipv4 multiplier multiplier	Specifies a number that is used as a multiplier with the
	Example:	minimum interval to determine BFD control packet failure detection times and transmission intervals for IPv4 BFD
	RP/0/RP0/CPU0:router(config-if)#bfd address-family ipv4 multiplier 30	sessions on bundle member links. The range is from 2 to 50. The default is 3.
		Note Although the command allows you to configure a minimum of 2, the supported minimum is 3.
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		 No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.

Configure BFD over Bundles IETF Mode Support on a Per Bundle Basis

To configure BFD over Bundles IETF mode support on a per bundle basis use these steps:

SUMMARY STEPS

- 1. configure
- 2. interface Bundle-Ether bundle-id
- **3**. bfd mode ietf
- 4. bfd address-family ipv4 fast-detect
- 5. Use the commit or end command.
- 6. show bundle bundle-ether bundle-id

	Command or Action	Purpose
Step 1	configure	Enters mode.
	Example:	
	RP/0/RP0/CPU0:router# configure	

	Command or Action	Purpose
Step 2	interface Bundle-Ether <i>bundle-id</i> Example:	Enters interface configuration mode for the specified bundle ID.
	RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1	
Step 3	bfd mode ietf	Enables IETF mode for BFD over bundle for the specified bundle.
	Example:	
	RP/0/RP0/CPU0:router(config-if)# bfd mode ietf	
Step 4	bfd address-family ipv4 fast-detect	Enables IPv4 BFD sessions on the specified bundle.
	Example:	
	<pre>RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 fast-detect</pre>	
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session.
		end —Prompts user to take one of these actions:
		• Yes — Saves configuration changes and exits the configuration session.
		• No —Exits the configuration session without committing the configuration changes.
		• Cancel —Remains in the configuration session, without committing the configuration changes.
Step 6	show bundle bundle-ether bundle-id	Displays the selected bundle mode.