



Configuring OSPF and RIP for VOSS

Release 8.0.50 (VOSS for XA1400 Series)
9035649-02 Rev AA
August 2019

© 2017-2019, Extreme Networks, Inc.
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: www.extremenetworks.com/support/policies/software-licensing

Contents

Chapter 1: About this Document	6
Purpose	6
Conventions	6
Text Conventions	6
Documentation and Training	8
Getting Help	9
Providing Feedback to Us	10
Chapter 2: New in this Document	11
Notice about Feature Support	11
Chapter 3: Routing fundamentals	12
Routing Protocols	12
Chapter 4: OSPF	13
OSPF fundamentals	13
OSPF overview	13
Dijkstras algorithm	14
Autonomous system and areas	14
OSPF neighbors	18
Router types	19
OSPF Interfaces	19
OSPF and IP	24
OSPF Packets	24
Intra-area Link-state Advertisements	25
ASE routes	25
OSPF virtual links	26
OSPF ASBRs	26
OSPF metrics	28
OSPF security mechanisms	28
OSPF and route redistribution	30
OSPF configuration considerations	30
OSPF host route advertisements and nonbackbone areas	31
OSPF with switch clustering	31
OSPF Graceful Restart	32
Open Shortest Path First guidelines	33
OSPF configuration using CLI	38
Configuring OSPF globally	38
Configuring OSPF for a Port or VLAN	41
Viewing OSPF errors on a port	44
Configuring OSPF areas on the router	45
Viewing the OSPF area information	47

Configuring OSPF aggregate area ranges on the router.....	47
Viewing the OSPF area range information.....	49
Enabling automatic virtual links.....	49
Configuring an OSPF area virtual interface.....	50
Configuring an OSPF area on a VLAN or port.....	52
Configuring an OSPF host route.....	55
Configuring OSPF NBMA neighbors.....	56
Enabling or disabling Helper mode for OSPFv2.....	57
Applying OSPF route acceptance policies.....	58
Viewing the OSPF configuration information.....	59
Viewing the OSPF link-state database.....	60
Viewing the OSPF external link-state database.....	62
Configuring route redistribution to OSPF.....	62
Viewing the OSPF redistribution configuration information.....	66
Configuring interVRF route redistribution for OSPF.....	67
Forcing shortest-path calculation updates.....	70
Viewing the OSPF default cost information.....	70
Viewing the OSPF interface statistics.....	71
Viewing the OSPF timer information.....	72
Viewing the OSPF neighbor information.....	73
Viewing the OSPF authentication information.....	73
Viewing the OSPF performance statistics.....	74
Viewing the OSPF virtual link information.....	75
Viewing the VRF configurations.....	76
Viewing the VRFIDS.....	77
OSPF configuration using EDM.....	78
Configuring OSPF globally.....	78
Enabling OSPF globally.....	80
Configuring global default metrics.....	80
Configuring an OSPF interface.....	81
Changing an OSPF non-passive interface type.....	84
Changing an OSPF passive interface type.....	84
Viewing the OSPF advanced interface.....	85
Configuring NBMA interface neighbors.....	86
Configuring OSPF interface metrics.....	87
Viewing all OSPF-enabled interfaces.....	88
Configuring OSPF on a port.....	88
Configuring OSPF on a VLAN.....	91
Viewing graphs for OSPF on a VLAN.....	94
Creating stubby or not-so-stubby OSPF areas.....	96
Configuring stub area metrics advertised by an ABR.....	98
Inserting OSPF area aggregate ranges.....	98
Enabling automatic virtual links.....	100

Configuring a manual virtual interface.....	100
Viewing virtual neighbors.....	101
Configuring host routes.....	102
Enabling ASBR status.....	103
Managing OSPF neighbors.....	104
Viewing the link-state database.....	104
Configuring interVRF route redistribution policies.....	105
Configuring route redistribution to OSPF.....	106
Viewing OSPF status.....	108
Forcing shortest-path calculation updates.....	110
Chapter 5: RIP	112
RIP fundamentals.....	112
Routing Information Protocol.....	112
RIP and route redistribution.....	113
RIP configuration using CLI.....	114
Configuring RIP globally.....	114
Configuring RIP on an interface.....	116
Configuring route redistribution to RIP.....	119
Configuring interVRF route redistribution for RIP.....	121
Forcing a RIP Update for a Port or VLAN.....	122
Viewing the RIP redistribution configuration information.....	123
RIP configuration using EDM.....	124
Configuring RIP globally.....	124
Viewing RIP status.....	125
Configuring RIP interface compatibility.....	126
Configuring RIP on an interface.....	128
Configuring RIP on a port.....	130
Configuring RIP on a VLAN.....	132
Configuring interVRF route redistribution policies.....	134
Configuring route redistribution to RIP.....	135
Glossary	137

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Extreme Networks VSP 4000 Series
- Extreme Networks VSP 7200 Series
- Extreme Networks VSP 7400 Series
- Extreme Networks VSP 8000 Series (includes VSP 8200 Series and VSP 8400 Series)
- Extreme Networks VSP 8600 Series
- Extreme Networks XA1400 Series

This document provides procedures and conceptual information that you can use to configure the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) on the VOSS switches. The router uses these protocols to determine the best routes for data forwarding.

For information about the Border Gateway Protocol, see [Configuring BGP Services for VOSS](#).

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons

Icon	Alerts you to...
 Important:	A situation that can cause serious inconvenience.
 Note:	Important features or instructions.
 Tip:	Helpful tips and notices for using the product.
 Danger:	Situations that will result in severe bodily injury; up to and including death.
 Warning:	Risk of severe personal injury or critical loss of data.
 Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command. If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code> , you can enter <code>cfm maintenance-domain maintenance-level 4</code>.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: - Click OK. - On the Tools menu, choose Options.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.
Brackets ([])	Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.

Table continues...

Convention	Description
Ellipses (...)	An ellipsis (...) indicates that you repeat the last element of the command as needed. For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>] ...</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.
<i>Italic Text</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: • <code>show ip route</code> • <code>Error: Invalid command syntax</code> <code>[Failed][2013-03-22 13:37:03.303</code> <code>-04:00]</code>
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the Configuration > Edit folders.
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code>, you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code>, but not both.

Documentation and Training

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation www.extremenetworks.com/documentation/

Table continues...

Archived Documentation (for earlier versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes
Hardware/Software Compatibility Matrices	https://www.extremenetworks.com/support/compatibility-matrices/
White papers, data sheets, case studies, and other product resources	https://www.extremenetworks.com/resources/

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For more information, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#)

For immediate support: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribing to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form with your information (all fields are required).
3. Select the products for which you would like to receive notifications.

 Note:

You can modify your product selections or unsubscribe at any time.

4. Click **Submit**.

Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team, you can do so in two ways:

- Use our short online feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Please provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this Document

There are no feature changes in this document.

Notice about Feature Support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not appear on your hardware, it is not supported.

For information about feature support, see [Release Notes for XA1400 Series](#).

For information about physical hardware restrictions, see your hardware documentation.

Chapter 3: Routing fundamentals

Use the information in this section to help you understand IP routing.

For more information about how to use the command line interface (CLI), see [Configuring User Interfaces and Operating Systems for VOSS](#).

Routing Protocols

Routers and routing switches use routing protocols to exchange reachability information. Routers use a routing protocol to advertise available paths on which the router can forward data. The routers use the protocol to determine the most efficient path to use. Routers use dynamic routing protocols to avoid sending data to inoperable links, and to send data to links that generally result in the fastest transmission times.

The switch routes frames using one of the following dynamic unicast IP routing protocols for path selection:

- Routing Information Protocol version 1 (RIPv1) (RFC 1058)
- RIPv2 (RFC 2453)
- Open Shortest Path First version 2 (OSPFv2) (RFC 2328)
- OSPFv3 (RFC 2740)
- Border Gateway Protocol version 4 (BGPv4) (RFC 1771)

Unlike static IP routing, where you must create a manual entry in the routing table to specify a routing path, dynamic IP routing uses a learning approach to determine the paths and routes to other routers. Dynamic routing uses two basic types of routing: distance vector and link-state. Routing Information Protocol (RIP) is a distance vector protocol and Open Shortest Path First (OSPF) is a link-state protocol.

The switch uses routing protocols like OSPF and RIP to populate routing tables. Routers use a routing protocol to exchange network topology information. A router uses the IP address of an incoming data packet to send the packet according to the routing tables.

The most commonly used unicast routing protocols include OSPF, RIP, and BGP. For more information about BGP, see [Configuring BGP Services for VOSS](#). For information about multicast routing protocols, see [Configuring IP Multicast Routing Protocols for VOSS](#). For information about OSPFv3 routing protocols, see [Configuring IPv6 Routing for VOSS](#).

Chapter 4: OSPF

This chapter provides concepts and configuration procedures for Open Shortest Path First (OSPF).

OSPF fundamentals

Use the information in these sections to help you understand Open Shortest Path First (OSPF).

OSPF is an Interior Gateway Protocol (IGP) that distributes routing information between routers that belong to a single autonomous system (AS). Intended for use in large networks, OSPF is a link-state protocol that supports IP subnets, Type of Service (TOS)-based routing, and tagging of externally-derived routing information.

For information about the Border Gateway Protocol (BGP), see [Configuring BGP Services for VOSS](#).

OSPF overview

In an OSPF network, each router maintains a link-state database that describes the topology of the AS. The database contains the local state for each router in the AS, including its usable interfaces and reachable neighbors. Each router periodically checks for changes in its local state and shares detected changes by flooding link-state advertisements (LSA) throughout the AS. Routers synchronize their topological databases based on the sharing of information from LSAs.

From the topological database, each router constructs a shortest-path tree, with itself as the root. The shortest-path tree provides the optimal route to each destination in the AS. Routing information from outside the AS appears on the tree as leaves.

OSPF routes IP traffic based on the destination IP address, subnet mask, and IP TOS.

In large networks, OSPF offers the following benefits:

- fast convergence
After network topology changes, OSPF recalculates routes quickly.
- minimal routing protocol traffic

Unlike distance vector routing protocols, such as Routing Information Protocol (RIP), OSPF generates a minimum of routing protocol traffic.

- load sharing

OSPF provides support for Equal Cost Multipath (ECMP) routing. If several equal-cost routes to a destination exist, ECMP distributes traffic equally among them.

- type of service

OSPF can calculate separate routes for each IP TOS.

Dijkstras algorithm

A separate copy of the OSPF routing algorithm (Dijkstra's algorithm) runs in each area. Routers that connect to multiple areas run multiple copies of the algorithm. The sequence of processes governed by the routing algorithm is as follows:

1. After a router starts, it initializes the OSPF data structures, and then waits for indications from lower-level protocols that the router interfaces are functional.
2. A router then uses the Hello protocol to discover neighbors. On point-to-point and broadcast networks the router dynamically detects neighbors by sending hello packets to the multicast address AllSPFRouters. On Non-Broadcast Multiple Access (NBMA) networks, you must provide some configuration information to discover neighbors.
3. On all multiaccess networks (broadcast or nonbroadcast), the Hello protocol elects a designated router (DR) for the network.
4. The router attempts to form adjacencies with some of its neighbors. On multiaccess networks, the DR determines which routers become adjacent. This behavior does not occur if you configure a router as a passive interface because passive interfaces do not form adjacencies.
5. Adjacent neighbors synchronize their topological databases.
6. The router periodically advertises its link state, and does so after its local state changes. LSAs include information about adjacencies, enabling quick detection of dead routers on the network.
7. LSAs flood throughout the area to ensure that all routers in an area have an identical topological database.
8. From this database each router calculates a shortest-path tree, with itself as the root. This shortest-path tree in turn yields a routing table for the protocol.

Autonomous system and areas

The AS subdivides into areas that group contiguous networks, routers that connect to these networks, and attached hosts. Each area has a topological database, which is invisible from outside the area. Routers within an area know nothing of the detailed topology of other areas. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS like a single link-state domain.

You can attach a router to more than one area. When you perform this action, you can maintain a separate topological database for each connected area. Two routers within the same area maintain

an identical topological database for that area. Each area uses a unique area ID and the area ID 0.0.0.0 is reserved for the backbone area.

The router routes packets in the AS based on their source and destination addresses. If the source and destination of a packet reside in the same area, the router uses intra-area routing. If the source and destination of a packet reside in different areas, the router uses inter-area routing. Intra-area routing protects the area from bad routing information because it does not use routing information obtained from outside the area. Inter-area routing must pass through the backbone area. For more information about the backbone area, see [Backbone area](#) on page 15.

In large networks with many routers and networks, the link-state database (LSDB) and routing table can become excessively large. Large route tables and LSDBs consume memory. The processing of link-state advertisements results in additional CPU cycles to make forwarding decisions. To reduce these undesired effects, you can divide an OSPF network into subdomains called areas.

An area comprises a number of OSPF routers that have the same area identification (ID).

By dividing a network into multiple areas, the router maintains a separate LSDB, which consists of router LSAs and network LSAs, for each area. Each router within an area maintains an LSDB only for the area to which it belongs. Area router LSAs and network LSAs do not flood beyond the area borders.

The impact of a topology change is localized to the area in which it occurs. The only exception is for the area border router (ABR), which must maintain an LSDB for each area to which they belong. The area border routers advertise changes in topology to the remainder of the network by advertising summary LSAs.

A 32-bit area ID, expressed in IP address format (x.x.x.x), identifies areas. Area 0 is the backbone area and distributes routing information to all other areas.

If you use multiple areas, they must all attach to the backbone through an ABR, which connects area 0.0.0.0 to the nonbackbone areas. If you cannot physically and directly connect an area through an ABR to area 0, you must configure a virtual link to logically connect the area to the backbone area.

Backbone area

The backbone area consists of the following network types:

- networks and attached routers that do not exist in other areas
- routers that belong to multiple areas

The backbone is usually contiguous but you can create a noncontiguous area by configuring virtual links.

You can configure virtual links between two backbone routers that have an interface to a nonbackbone area. Virtual links belong to the backbone and use intra-area routing only.

The backbone distributes routing information between areas. The topology of the backbone area is invisible to other areas, while it knows nothing of the topology of those areas.

In inter-area routing, a packet travels along three contiguous paths in a point-to-multipoint configuration:

- an intra-area path from the source to an ABR
- a backbone path between the source and destination areas

- another intra-area path to the destination

The OSPF routing algorithm finds the set of paths that has the smallest cost. The topology of the backbone dictates the backbone paths used between areas. OSPF selects inter-area paths by examining the routing table summaries for each connected ABR. The router cannot learn OSPF routes through an ABR unless it connects to the backbone or through a virtual link.

Stub area

Configure a stub area at the edge of the OSPF routing domain. A stub area has only one ABR. A stub area does not receive LSAs for routes outside its area, which reduces the size of its link-state database. A packet destined outside the stub area is routed to the ABR, which examines it before forwarding the packet to the destination. The network behind a passive interface is treated as a stub area and does not form adjacencies. The network is advertised into the OSPF area as an internal route.

Not so stubby area

A not-so-stubby area (NSSA) prevents the flooding of external LSAs into the area by replacing them with a default route. An NSSA can import small stub (non-OSPF) routing domains into OSPF. Like stub areas, NSSAs are at the edge of an OSPF routing domain. Non-OSPF routing domains attach to the NSSAs to form NSSA transit areas. Accessing the addressing scheme of small stub domains permits the NSSA border router to also perform manual aggregation.

In an OSPF NSSA, the NSSA N/P bit notifies the ABR which external routes to advertise to other areas. If the NSSA N/P bit is set (the value is 1), the ABR exports the external route. This configuration is the default. When the NSSA N/P bit is not set (the value is 0), the ABR drops the external route. You can create a route policy to manipulate the N/P bit.

Multiarea OSPF configuration

The following figure shows five devices (R1 to R5) in a multi-area configuration.




Figure 1: Multiarea configuration example

The following list explains the configuration for devices R1 through R5:

- R1 is an OSPF AS boundary router (ASBR) that is associated with OSPF Area 0 and OSPF Area 3. R1 distributes a default route for Internet traffic.
- R2 is an OSPF stub ABR for OSPF Area 2 and ABR to OSPF Area 3.
- R3 is an OSPF ASBR and distributes OSPF to RIP and RIP to OSPF.
- R4 is an OSPF internal router in Area 3.
- R5 is an internal OSPF subrouter in Area 2.
- All OSPF interfaces are brouter ports except R5.

Network 203.0.113.0/24 on R5 uses a VLAN configuration instead of a brouter port. This example uses brouter ports rather than VLANs because the spanning tree algorithm is disabled by default if you use brouter interfaces.

- All interfaces are Ethernet; therefore, the OSPF interfaces are broadcast, except the circuitless IP (CLIP) interfaces, which are passive.
- The interface priority on R5 is 0; therefore, R5 cannot become a DR.
- Configure the OSPF router priority so that R1 becomes the DR (priority 100) and R2 becomes the backup designated router (BDR) with a priority value of 50.

Use stub or NSSA areas to reduce the LSDB size by excluding external LSAs. The stub ABR advertises a default route into the stub area for all external routes.

OSPF neighbors

In an OSPF network, two routers that have an interface to the same network are neighbors. Routers use the Hello protocol to discover their neighbors and to maintain neighbor relationships. On a broadcast or point-to-point network, the Hello protocol dynamically discovers neighbors. On an NBMA network, you must manually configure neighbors for the network.

The Hello protocol provides bidirectional communication between neighbors. Periodically, OSPF routers send hello packets over all interfaces. Included in these hello packets is the following information:

- router priority
- router hello timer and dead timer values
- list of routers that sent the router hello packet on this interface
- router choice for DR and backup designated router (BDR)

Bidirectional communication is determined after one router discovers itself listed in the hello packet of its neighbor.

NBMA interfaces whose router priority is a positive, nonzero value are eligible to become DRs for the NBMA network and are configured with a list of all attached routers. The neighbors list includes each neighbor IP address and router priority. In an NBMA network, a router with a priority other than zero is eligible to become the DR for the NBMA network. You must manually configure the IP address, mask, and router priority of neighbors on routers that are eligible to become the DR or BDR for the network.

Log messages indicate when an OSPF neighbor state change occurs. Each log message indicates the previous state and the new state of the OSPF neighbor. The log message generated for system traps also indicates the previous state and the current state of the OSPF neighbor.

Neighbors can form an adjacency to exchange routing information. After two routers form an adjacency, they perform a database exchange process to synchronize their topological databases. After the databases synchronize, the routers are fully adjacent. Adjacency conserves bandwidth because, from this point, the adjacent routers pass only routing change information.

All routers connected by a point-to-point network or a virtual link always form an adjacency. All routers on a broadcast or NBMA network form an adjacency with the DR and the BDR.

In an NBMA network, before the routers elect a DR, the router sends hello packets only to those neighbors eligible to become a DR. The NBMA DR forms adjacencies only with its configured neighbors and drops all packets from other sources. The neighbor configuration also notifies the router of the expected hello behavior for each neighbor.

If a router receives a hello packet from a neighbor with a priority different from that which is already configured for the neighbor, the router can automatically change the configured priority to match the dynamically learned priority.

Router types

To limit the amount of routing protocol traffic, the Hello protocol elects a DR and a BDR on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information, which on a large network can mean significant routing protocol traffic, all routers on the network form adjacencies with the DR and the BDR only, and send link-state information to them. The DR redistributes this information to every other adjacent router.

If the BDR operates in backup mode, it receives link-state information from all routers on the network and listens for acknowledgements. If the DR fails, the BDR can transition quickly to the role of DR because its routing tables are up-to-date.

Routers in an OSPF network can have various roles depending on how you configure them. The following table describes the router types you can configure in an OSPF network.

Table 3: Router types in an OSPF network

Router type	Description
AS boundary router	A router that attaches at the edge of an OSPF network is an ASBR. An ASBR generally has one or more interfaces that run an interdomain routing protocol such as Border Gateway Protocol. In addition, a router that distributes static routes or RIP routes into OSPF is an ASBR. The ASBR forwards external routes into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area border router	A router that attaches to two or more areas inside an OSPF network is an ABR. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is an IR. Unlike ABRs, IRs have topological information only about the area in which they reside.
Designated router	In a broadcast or NBMA network, the routers elect a single router as the DR for that network. A DR makes sure that all routers on the network synchronize and advertises the network to the rest of the AS.
Backup designated router	A BDR is elected in addition to the DR and, if the DR fails, can assume the DR role quickly.

OSPF Interfaces

Configure an OSPF interface, or link, on an IP interface. An IP interface can be either a single link (brouter port) or a logical interface configured on a VLAN (multiple ports). The state information associated with the interface is obtained from the underlying lower-level protocols and the routing protocol itself.

! Important:

To change the interface type of an enabled OSPF interface, you must first disable it, change the type, and then reenabling it. For an NBMA interface, you must first delete manually configured neighbors.

OSPF network types allow OSPF-neighboring between routers over various types of network infrastructures. You can configure each interface to support various network types. The following table describes the supported OSPF network interface types:

Table 4: OSPF Network Types

Network interface type	Description
Broadcast Interfaces on page 20	Broadcast interfaces automatically discover every OSPF router on the network by sending OSPF hello packets to the multicast group AllSPFRouters (224.0.0.5). Neighboring is automatic and requires no configuration.
Non-Broadcast Multiple Access Interfaces on page 20	The NBMA network type models network environments that do not have native Layer 2 broadcast or multicast capabilities, such as Frame Relay and X.25. OSPF hello packets are unicast to manually configured neighbors.
Passive Interfaces on page 24	A passive interface is an interfacing network in OSPF that does not generate LSAs or form adjacencies. Use a passive interface on an access network or on an interface used for BGP peering. Using passive interfaces limits the amount of CPU cycles required to perform the OSPF routing algorithm.

Broadcast Interfaces

Broadcast interfaces support many attached routers and can address a single physical message to all attached broadcast routers (sent to AllSPFRouters and AllDRouters).

Broadcast interfaces dynamically discover neighboring routers using the OSPF Hello protocol. Each pair of routers on a broadcast network, such as an Ethernet, communicate directly.

Non-Broadcast Multiple Access Interfaces

An NBMA network interconnects multiple devices through point-to-point links. NBMA does not use broadcast and multicast data transmission.

NBMA interfaces support many routers, but cannot broadcast. NBMA networks perform the following activities:

- statically establish OSPF neighbor relationships
You must establish neighbor relationships because hub-and-spoke Wide Area Network (WAN) topologies do not support any-to-any broadcasting.
- control meshed WAN connections

In contrast to a broadcast network, where some OSPF protocol packets are multicast (sent to AllSPFRouters and AllDRouters), OSPF packets on an NBMA interface are replicated and sent in turn to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination address AllSPFRouters and AllDRouters.

The following figure shows an example of four routers attached to an NBMA subnet. The NBMA segment uses a single IP subnet and each router uses an IP address within the subnet.




Figure 2: NBMA Subnet

NBMA Interface Operations and Parameters

OSPF treats an NBMA network much like it treats a broadcast network. Because many routers attach to the network, the Hello protocol elects a DR to generate the network link-state advertisements.

Because the NBMA network does not broadcast, you must manually configure neighbors for each router eligible to become DR (those networks with a positive, nonzero router priority value). You must also configure a poll interval for the network.

NBMA interfaces with a positive, nonzero router priority can become DR for the NBMA network and contain a list of all attached routers, or neighbors. This neighbors list includes each neighbor IP address and router priority.

The router uses neighbor information both during and after the DR election process. After an interface to a nonbroadcast network with a nonzero priority initializes, and before the Hello protocol elects a DR, the router sends hello packets only to those neighbors eligible to become DR. After the Hello protocol elects a DR, it forms adjacencies only with its configured neighbors and drops all packets from other sources. This neighbor configuration also notifies the router of the expected hello behavior of each neighbor.

If a router eligible to become the DR receives a hello packet from a neighbor that shows a different priority from that which is already configured for this neighbor, the DR changes the configured priority to match the dynamically learned priority.

Configure an NBMA interface with a poll interval. The poll interval designates the interval at which the router sends hello packets to inactive neighboring routers. The router typically sends hello packets at the Hello interval, for example, every 10 seconds. If a neighboring router becomes inactive, or if the router does not receive hello packets for the established RouterDeadInterval period, the router sends hello packets at the specified poll interval, for example, every 120 seconds.

You must configure a neighbors list for the DR to allow an NBMA network to send hello packets. If the router is eligible to become a DR, it periodically sends hello packets to all neighbors that are also eligible. The effect of this action is that two eligible routers always exchange hello packets, which is necessary for the correct DR election. You can minimize the number of hello packets by minimizing the number of eligible routers on a nonbroadcast network.

After the Hello protocol elects a DR, it sends hello packets to all manually configured neighbors to synchronize their link-state databases, establish itself as the DR, and identify the BDR.

If a router is not eligible to become DR, it periodically sends hello packets to both the DR and the BDR. The router also sends a hello packet in reply to a hello packet received from an eligible neighbor (other than the current DR and BDR). This process establishes an initial bidirectional relationship with a potential DR.

When a router sends hello packets to a neighbor, the neighbor state determines the interval between hello packets. If the neighbor is in the down state, the router sends hello packets at the designated poll interval, for example, every 120 seconds. Otherwise, the router sends hello packets at the designated hello interval, for example, every 10 seconds.

OSPF and NBMA Example: Adjacency Formation

In an NBMA network, as in a broadcast network, all routers become adjacent to the DR and the BDR. The adjacencies form after you assign the router priorities, configure the neighbors, and the Hello protocol elects the network DR.

The following figure shows an NBMA subnet with router priorities and manually configured neighbors.




Figure 3: NBMA Subnet Configuration Example

Because R1 and R2 have a router priority of 0, they are not eligible to become the DR. Also, R1 and R2 do not require configuration of a neighbors list; R1 and R2 discover neighbors dynamically through the Hello protocol.

R3 and R4 both have a positive, nonzero priority and are eligible to become the DR. Manually configure neighbor lists on R3 and R4.

To create this NBMA network, configure the following parameters:

1. On each router: NBMA interface type, poll interval, router priority
2. On R3: R1, R2, and R4 as neighbors
3. On R4: R1, R2, and R3 as neighbors

If all routers start at the same time, the routers perform the following steps:

1. R3 and R4 send each other a hello packet to elect a DR.
2. The Hello protocol elects R3 as the DR, and R4 as the BDR.
3. R3 (DR) and R4 (BDR) send hello packets to all other routers on the NBMA subnet to synchronize their link-state databases and establish themselves as DR and BDR.

4. R1 and R2 reply to R3 and R4.
5. R3 and R4 each form three adjacencies (one with each router on the NBMA subnet).
6. R1 and R2 each form two adjacencies (one with the DR and one with the BDR).

Passive Interfaces

Use a passive interface to enable an interface to advertise into an OSPF domain while limiting its adjacencies.

After you change the interface type to passive, the router advertises the interface into the OSPF domain as an internal stub network with the following behaviors:

- does not send hello packets to the OSPF domain
- does not receive hello packets from the OSPF domain
- does not form adjacencies in the OSPF domain

If you configure an interface as passive, the router advertises it as an OSPF internal route. If the interface is not a passive interface, to advertise a network into OSPF and not form OSPF adjacencies, you must configure the interface as nonOSPF, and the router must redistribute the local network as an autonomous system external (ASE) LSA.

OSPF and IP

OSPF runs over IP, which means that an OSPF packet transmits with an IP data packet header. The protocol field in the IP header is 89, which identifies it as an OSPF packet and distinguishes it from other packets that use an IP header.

An OSPF route advertisement expresses a destination as an IP address and a variable-length mask. Together, the address and the mask indicate the range of destinations to which the advertisement applies.

Because OSPF can specify a range of networks, it can send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 192.0.2.0 with a mask of 255.255.0.0 describes a single route to destinations 192.0.2.0 to 192.0.2.255.

OSPF Packets

All OSPF packets start with a 24-octet header that contains information about the OSPF version, the packet type and length, the ID of the router that transmits the packet, and the ID of the OSPF area that sends the packet. An OSPF packet is one of the following types:

- The router transmitted hello packets between neighbors and never forwards them. The Hello protocol requires routers to send hello packets to neighbors at predefined hello intervals. A neighbor router that does not receive a hello packet declares the other router dead.
- The router exchanges DD packets after neighboring routers establish a link, which synchronizes their LSDBs.

- Link-state request packets describe one or more link-state advertisements that a router requests from its neighbor. Routers send link-state requests if the information received in DD packets from a neighbor is not consistent with its own link-state database.
- Link-state update packets contain one or more LSAs and the router sends them following a change in network conditions.
- The router sends link-state acknowledgement packets to acknowledge receipt of link-state updates. Link-state acknowledgement packets contain the headers of the received LSAs.

Intra-area Link-state Advertisements

OSPF does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information in the form of LSAs throughout the area or AS. LSAs in OSPF are one of the following five types:

- A router links advertisement is flooded only within the area and contains information about neighbor routers and the LANs to which the router attaches. A backbone router can flood router link advertisements within the backbone area.
- A DR on a LAN generates network links advertisement to list all routers on that LAN, and floods network links advertisements only within the area. A backbone DR can flood network links advertisements within the backbone area.
- An ABR floods a network summary link advertisement into an area and describes networks that are reachable outside the area. An ABR attached to two areas generates a different network summary link advertisement for each area. ABRs also generate area summary link advertisements that contain information about destinations within an area that are flooded to the backbone area.
- An ASBR summary link advertisement describes the cost of the path to an ASBR from the router that generates the advertisement.
- An ASBR sends an ASE link advertisement to describe the cost of the path to a destination outside the AS from the ASBR that generates the advertisement. This information is flooded to all routers in the AS.

ASE routes

OSPF considers the following routes as ASE routes:

- a route to a destination outside the AS
- a static route
- a default route
- a route derived by RIP
- a directly connected network that does not run OSPF

OSPF virtual links

On an OSPF network, a switch that acts as an ABR must connect directly to the backbone. If no physical connection is available, you can automatically or manually configure a virtual link.

An automatic virtual link can provide redundancy support for critical network connections. Automatic virtual linking creates virtual paths for vital traffic paths in your OSPF network. If a connection fails on the network, such as after an interface cable that provides connection to the backbone (either directly or indirectly) disconnects from the switch, the virtual link is available to maintain connectivity.

Use automatic virtual linking to ensure that a link is created to another router. If automatic virtual linking uses more resources than you want to expend, creating a manual virtual link can be the better solution. Use this approach to conserve resources and control virtual links in the OSPF configuration.

On the switch, OSPF behavior follows OSPF standards; the router cannot learn OSPF routes through an ABR unless the ABR connects to the backbone or through a virtual link.

The following figure shows how to configure a virtual link between the ABR in area 2.2.2.2 and the ABR in area 0.0.0.0.




Figure 4: Virtual link between ABRs through a transit area

To configure a virtual link between the ABRs in area 1 and area 3, define area 2 as the transit area between the other two areas, and identify R2 as the neighbor router through which R2 must send information to reach the backbone through R1.

OSPF ASBRs

ASBRs advertise nonOSPF routes into OSPF domains so that they can pass through the OSPF routing domain. A router can function as an ASBR if one or more interfaces connects to a nonOSPF network, for example, RIP, BGP, or Exterior Gateway Protocol (EGP).

An ASBR imports external routes into the OSPF domain by using ASE LSAs (LSA type 5) originated by the ASBR.

ASE LSAs flood across area borders. When an ASBR imports external routes, it imports OSPF route information using external type 1 or type 2 metrics. The result is a four-level routing hierarchy, as shown in the following table, according to routing preference.

Table 5: ASBR routing hierarchy

Level	Description
1	Intra-area routing
2	Inter-area routing
3	External type 1 metrics
4	External type 2 metrics

The use of these metrics results in a routing preference from most preferred to least preferred of

- routing within an OSPF area
- routing within the OSPF domain
- routing within the OSPF domain and external routes with external type 1 metrics
- routing within the OSPF domain and external routes with external type 2 metrics

For example, an ASBR can import RIP routes into OSPF with external type 1 metrics. Another ASBR can import Internet routes and advertise a default route with an external type 2 metric. This results in RIP-imported routes that have a higher preference than the Internet-imported default routes. In reality, BGP Internet routes must use external type 2 metrics, whereas RIP imported routes must use external type 1 metrics.

Routes imported into OSPF as external type 1 are from IGPs whose external metric is comparable to OSPF metrics. With external type 1 metrics, OSPF adds the internal cost of the ASBR to the external metric. EGPs, whose metric is not comparable to OSPF metrics, use external type 2 metrics. External type 2 metrics use only the internal OSPF cost to the ASBR in the routing decision.

To conserve resources, you can limit the number of ASBRs in your network or specifically control which routers perform as ASBRs to control traffic flow.

Area link-state advertisements

The following table explains the seven LSA types exchanged between areas. LSAs share link-state information among routers. LSAs typically contain information about the router and its neighbors. OSPF generates LSAs periodically to ensure connectivity or after a change in state of a router or link (that is, up or down).

Table 6: OSPF LSA types

LSA type	Description	Area of distribution
1	A router originates type 1 LSAs (router LSAs) to describe its set of active interfaces and neighbors.	Passed only within the same area
2	Type 2 LSAs (network LSAs) describe a network segment such as broadcast or NBMA. In a broadcast network, the DR originates network LSAs.	Passed only within the same area

Table continues...

LSA type	Description	Area of distribution
3	The ABR originates type 3 LSAs (network-summary LSAs) to describe the networks within an area.	Passed between areas
4	Type 4 LSAs (ASBR-summary LSAs) advertise the location of the ASBRs from area to area.	Passed between areas
5	Type 5 LSAs (ASE LSAs) describe networks outside of the OSPF domain. The ASBR originates type 5 LSAs. In stub and NSSA areas, a single default route replaces type 5 LSA routes.	Passed between areas
6	Type 6 LSAs (group membership LSAs) identify the location of multicast group members in multicast OSPF.	Passed between areas
7	Type 7 LSAs import external routes in OSPF NSSAs.	Translated between areas

OSPF metrics

For OSPF, the best path to a destination is the path that offers the least-cost metric (least-cost delay). You can configure OSPF cost metrics to specify preferred paths. You can configure metric speed globally or for specific ports and interfaces on the network. In addition, you can control redistribution options between nonOSPF interfaces and OSPF interfaces.

Assign default metric speeds for different port types, such as 10 Mb/s or 1 Mb/s ports. You can specify a new metric speed for an IP interface. An IP interface can be a brouter port or a VLAN.

RFC1583 states the following:

"OSPF supports two types of external metrics. Type 1 external metrics are equivalent to the link state metric. Type 2 external metrics are greater than the cost of path internal to the Autonomous System. Use of Type 2 external metrics assumes that routing between Autonomous Systems is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link state metrics."

"Both Type 1 and Type 2 external metrics can be present in the Autonomous System at the same time. In that event, Type 1 external metrics always take precedence."

OSPF security mechanisms

The switch implementation of OSPF includes security mechanisms to prevent unauthorized routers from attacking the OSPF routing domain. These security mechanisms prevent a malicious person from joining an OSPF domain and advertising false information in the OSPF LSAs. Likewise, security prevents a misconfigured router from joining an OSPF domain.

Simple password

The simple password security mechanism is a simple-text password; only routers that contain the same authentication ID in their LSA headers can communicate with each other.

It is recommended that you do not use this security mechanism because the system stores the password in plain text. A user or system can read the password from the configuration file or from the LSA packet.

Message Digest 5

Message Digest 5 (MD5) for OSPF security provides standards-based (RFC1321) authentication using 128-bit encryption, usually expressed as a 32-digit hexadecimal number. When you use MD5 for OSPF security, it is almost impossible for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

If you use MD5, each OSPF packet has a message digest appended to it. The digest must match between the sending and receiving routers. Both the sending and receiving routers calculate the message digest based on the MD5 key and padding, and then compare the results. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet.

Secure hash algorithm 1

The secure hash algorithm 1 (SHA-1) is a cryptographic hash function that uses 160-bit encryption, usually given in a 40 digit hexadecimal number. SHA-1 is one of the most widely used of the existing SHA hash functions and is more secure than MD5.

SHA-1 takes a variable length input message and SHA-1 creates a fixed length output message referred to as the hash, or message digest, of the original message. If you use SHA-1 with OSPF, each OSPF packet has a message digest appended to it.

The message digest or hash must match between the sending and receiving routers. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet. The hash functions produce a type of checksum or summary of the input.

It is almost impossible to determine the original input message based on the output hash message.

A cryptographic hash function is fully defined and uses no secret key.

Secure hash algorithm 2

Secure hash algorithm 2 (SHA-2) is also a cryptographic hash function. SHA-2 updates SHA-1 and offers six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits message digest size values. Output size depends on the hash function, so, for instance, SHA-256 is 256 bits.

SHA-2 is more secure than SHA-1 and MD5.

SHA-2 works similarly to SHA-1, in that SHA-2 takes a variable length input message and creates a fixed length output message referred to as the hash, or message digest, of the original message. If you use SHA-2 with OSPF, each OSPF packet has a message digest appended to it. Among the differences in SHA-2 from SHA-1 are an increased bit encryption length.

Similarly with other hash functions, for SHA-2, the message digest or hash must match between the sending and receiving routers. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet. The hash functions produce a type of checksum or summary of the input.

OSPF and route redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if OSPF routes exist in a router and they must travel through a BGP network, then configure redistribution of OSPF routes through BGP. This configuration sends OSPF routes to a router that uses BGP.

You can redistribute routes

- on an interface basis
- on a global basis between protocols on a single VRF instance (intraVRF)
- between the same or different protocols on different VRF instances (interVRF)

To configure interface-based redistribution, configure a route policy, and then apply it to the interface. Configure the match parameter to the protocol from which to learn routes.

Use the `ip ospf redistribute` command to accomplish the (intraVRF) redistribution of routes through OSPF, so that OSPF redistribution occurs globally on all OSPF-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

OSPF route redistribution and DvR

DvR Controllers redistribute routes (direct routes, static routes and the default route) into the DvR domain. You can configure redistribution of DvR host routes into OSPF.

For information on DvR, see [Configuring IPv4 Routing for VOSS](#).

OSPF configuration considerations

This section describes considerations to keep in mind as you configure OSPF.

OSPF host route advertisements and nonbackbone areas

The switch does not associate a host route with a specific area. Therefore, if you create a host route in a nonbackbone area, nonbackbone (nonOSPF core) areas do not advertise it.

For example, in an OSPF network with multiple areas, including areas not adjacent to the core, which use virtual links, a host route on a router that belongs to a nonOSPF core area is not advertised on noncore routers.

To ensure host route advertisement, disable and enable OSPF on the noncore routers.

OSPF with switch clustering

If the network loses the DR, the BDR immediately becomes the new DR on the broadcast segment. After OSPF elects the new DR, all routers perform an SPF run and issue new LSAs for the segment. The new DR generates a new network LSA for the segment and every router on the segment must refresh the router LSA.

Each router performs the SPF run as soon as it detects a new DR. Depending on the speed of the router, the router can perform the SPF run before it receives the new LSAs for the segments, which requires a second SPF run to update and continue routing across the segment. The OSPF hold-down timer does not permit two consecutive SPF runs within the value of the timer. This limitation can lead to traffic interruption of up to 10 seconds.

In a classical OSPF routed design, this situation never causes a problem because OSPF runs over multiple segments so even if a segment is not usable, routes are recalculated over alternative segments. Typical Routed Split MultiLink Trunking (RSMLT) designs only deploy a single OSPF routed vlan, which constitutes a single segment.

You can use RSMLT in a configuration with dual core VLANs to minimize traffic interruption when the network loses the DR. This configuration creates a second OSPF core VLAN, forcing different nodes to become the DR for each VLAN. Each OSPF core VLAN has a DR (priority of 100) and no BDRs. This configuration does not require a BDR because the two VLANs provide backup for each other from a routing perspective. See the following figure for a network example.




Figure 5: RSMLT with dual core VLANs

OSPF Graceful Restart

In many OSPF networks, OSPF routers remove a restarting OSPF router from the network topology, if the router is restarted. This action causes all OSPF routers to re-converge and route around the restarting router. The OSPF Graceful Restart feature is an OSPF enhancement to allow an OSPF router to stay on the forwarding path when the software is restarting.

This feature is documented under RFC 3623 for OSPFv2 (IPv4) and RFC 5187 for OSPFv3 (IPv6). The switch software supports only helper mode for both OSPFv2 and OSPFv3 protocols.

Helper Mode

Helper mode is a part of the OSPF Graceful restart feature. Helper mode uses the OSPF routers to help other OSPF routers on the network stay on the forwarding path while the software is restarting. The OSPF router sends a type of LSA called a GRACE-LSA to inform the other OSPF routers that it is restarting the software. When an OSPF router receives a GRACE-LSA from a neighbor OSPF Router, it enters the Helper mode for that neighbor on that network. An OSPF router supports Helper mode by default.

Operations of Helper Mode

The following section describes the operations in the Helper mode:

- Entering Helper mode — An OSPF router enters the Helper mode provided the following conditions are true:
 - The router is fully adjacent with the neighbor already.
 - No changes have been made in the LSDB since the neighbor router started.

- The grace period has not expired.
- Local policy configured parameters allow it to help the neighbor.
- The router is not in the process of restarting itself.

The OSPF router will not help the neighbor if any of the above conditions are not met.

If the OSPF router is already helping a neighbor, and receives another GRACE-LSA from the neighbor, it accepts the latest GRACE-LSA, and updates the grace period accordingly. The OSPF router in Helper mode continues to advertise its LSAs like the neighbor it is helping is still full, until any changes are made on the network during the grace period.

- Exiting Helper mode — An OSPF router exits the Helper mode, under the following conditions:
 - The GRACE-LSA is flushed. It means graceful restart has successfully terminated.
 - The GRACE-LSA's grace period expires.
 - There is a network topology change.

When an OSPF router exits Helper mode, the following actions occur:

- It recalculates the DR for the network.
- It re-originates its router LSA.
- If it is the DR, it re-originates the network LSA for the network.
- If it is a virtual link, it re-originates the router LSA for the virtual link transit area.

Open Shortest Path First guidelines

Use OSPF to ensure that the switch can communicate with other OSPF routers. This section describes some general design considerations and presents a number of design scenarios for OSPF.

OSPF LSA limits

To determine OSPF link-state advertisement (LSA) limits:

1. Use the command `show ip ospf area` to determine the LSA_CNT and to obtain the number of LSAs for a given area.
2. Use the following formula to determine the number of areas. Ensure the total is less than 16,000 (16K):

$$\sum Adj_N * LSA_CNT_N < 16k$$

N = 1 to the number of areas for each switch

Adj_N = number of adjacencies for each Area N

LSA_CNT_N = number of LSAs for each Area N

For example, assume that a switch has a configuration of three areas with a total of 18 adjacencies and 1000 routes. This includes:

- 3 adjacencies with an LSA_CNT of 500 (Area 1)

- 10 adjacencies with an LSA_CNT of 1000 (Area 2)
- 5 adjacencies with an LSA_CNT of 200 (Area 3)

Calculate the number as follows:

$$3*500+10*1000+5*200=12.5K < 16K$$

This configuration ensures that the switch operates within accepted scalability limits.

OSPF design guidelines

Follow these additional OSPF guidelines:

- OSPF timers must be consistent across the entire network.
- Use OSPF area summarization to reduce routing table sizes.
- Use OSPF passive interfaces to reduce the number of active neighbor adjacencies.
- Use OSPF active interfaces only on intended route paths.

Configure wiring-closet subnets as OSPF passive interfaces unless they form a legitimate routing path for other routes.

- Minimize the number of OSPF areas for each switch to avoid excessive shortest-path calculations.

The switch executes the Dijkstra algorithm for each area separately.

- Ensure that the OSPF dead interval is at least four times the OSPF hello-interval.
- Use MD5 authentication on untrusted OSPF links.
- Use stub or NSSAs as much as possible to reduce CPU overhead.

OSPF and CPU utilization

After you create an OSPF area route summary on an area border router, the summary route can attract traffic to the area border router for which the router does not have a specific destination route. Enabling ICMP unreachable-message generation on the switch can result in a high CPU utilization rate.

To avoid high CPU utilization, it is recommended that you use a black-hole static route configuration. The black-hole static route is a route (equal to the OSPF summary route) with a next hop of 255.255.255.255. This configuration ensures that all traffic that does not have a specific next-hop destination route is dropped.

OSPF network design examples

You can use OSPF routing in the core of a network.

The following figure describes a simple implementation of an OSPF network: enabling OSPF on two switches (S1 and S2) that are in the same subnet in one OSPF area.




Figure 6: Example 1: OSPF on one subnet in one area

The routers in the preceding figure use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.0.2.1.
- S2 has an OSPF router ID of 1.1.1.2, and the OSPF port uses an IP address of 192.0.2.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Enable IP forwarding on the switch.
3. Configure the IP address, subnet mask, and VLAN ID for the port.
4. Disable RIP on the port, if you do not need it.
5. Enable OSPF for the port.

After you configure S2, the two switches elect a designated router and a backup designated router. They exchange hello packets to synchronize their link state databases.

The following figure shows a configuration in which OSPF operates on three switches. OSPF performs routing on two subnets in one OSPF area. In this example, S1 directly connects to S2, and S3 directly connects to S2, but traffic between S1 and S3 is indirect, and passes through S2.




Figure 7: Example 2: OSPF on two subnets in one area

The routers in example 2 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.0.2.1.
- S2 has an OSPF router ID of 1.1.1.2, and two OSPF ports use IP addresses of 192.0.2.2 and 198.51.100.1.
- S3 has an OSPF router ID of 1.1.1.3, and the OSPF port uses an IP address of 198.51.100.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Insert IP addresses, subnet masks, and VLAN IDs for the OSPF ports on S1 and S3, and for the two OSPF ports on S2. The two ports on S2 enable routing and establish the IP addresses related to the two networks.
3. Enable OSPF for each OSPF port allocated with an IP address.

After you configure all three switches for OSPF, they elect a designated router and a backup designated router for each subnet and exchange hello packets to synchronize their link-state databases.

The following figure shows an example where OSPF operates on two subnets in two OSPF areas. S2 becomes the area border router for both networks.




Figure 8: Example 3: OSPF on two subnets in two areas

The routers in scenario 3 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1. The OSPF port uses an IP address of 192.0.2.1, which is in OSPF area 1.
- S2 has an OSPF router ID of 1.1.1.2. One port uses an IP address of 192.0.2.2, which is in OSPF area 1. The second OSPF port on S2 uses an IP address of 198.51.100.1, which is in OSPF area 2.
- S3 has an OSPF router ID of 1.1.1.3. The OSPF port uses an IP address of 198.51.100.2, which is in OSPF area 2.

The general method to configure OSPF for this three-switch network is:

1. On all three switches, enable OSPF globally.
2. Configure OSPF on one network.

On S1, insert the IP address, subnet mask, and VLAN ID for the OSPF port. Enable OSPF on the port. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 1, and enable OSPF on the port. Both routable ports belong to the same network. Therefore, by default, both ports are in the same area.

3. Configure three OSPF areas for the network.
4. Configure OSPF on two additional ports in a second subnet.

Configure additional ports and verify that IP forwarding is enabled for each switch to ensure that routing can occur. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 2, and enable OSPF on the port. On S3, insert the IP address, subnet mask, and VLAN ID for the OSPF port, and enable OSPF on the port.

The three switches exchange hello packets.

In an environment with a mix of switches and routers from different vendors, you may need to manually modify the OSPF parameter RtrDeadInterval to 40 seconds.

OSPF configuration using CLI

Configure Open Shortest Path First (OSPF) so that the switch can use OSPF routing to communicate with other OSPF routers and to participate in OSPF routing.

Configuring OSPF globally

Configure OSPF parameters on the switch so that you can control OSPF behavior on the system. The switch uses global parameters to communicate with other OSPF routers. Globally configure OSPF before you configure OSPF for an interface, port, or VLAN.

Before you begin

- Ensure that the switch has an IP interface.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf` to commands. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Configure the OSPF router ID:

```
router-id {A.B.C.D}
```

3. Configure the router as an autonomous system boundary router (ASBR):

```
as-boundary-router enable
```

Note:

Configure the following steps as and when needed.

4. Enable the automatic creation of OSPF virtual links:

```
auto-vlink
```

5. Configure the OSPF default metrics:

```
default-cost {ethernet|fast-ethernet|forty-gig-ethernet|gig-
ethernet|hundred-gig-ethernet|ten-gig-ethernet|twentyfive-gig-
ethernet} <1-65535>]
default-cost vlan <1-65535>]
```

6. Configure the OSPF hold-down timer value:

```
timers basic holddown <3-60>
```

7. Enable the RFC1583 compatibility mode:

```
rfc1583-compatibility enable
```

8. Enable the router to issue OSPF traps:

```
trap enable
```

9. Verify the OSPF configuration:

```
show ip ospf [vrf WORD<1-16>] [vrfsids WORD<0-512>]
```

10. Exit OSPF Router Configuration mode:

```
exit
```

You return to Global Configuration mode.

11. Enable OSPF for the switch:

```
router ospf enable
```

Example

Configure the OSPF router ID to 192.0.2.2, enable the automatic creation of OSPF virtual links, and enable traps. Configure the default cost metric for Ethernet to 101, for fast Ethernet to 110, and for gig-Ethernet, ten-gig-Ethernet, twentyfive-gig-ethernet, forty-gig-Ethernet, and hundred-gig-ethernet to 20, and vlan to 1. Configure the basic holdown to 10. Enable OSPF for the switch, and review the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#router-id 192.0.2.2
Switch:1(config-ospf)#auto-vlink
Switch:1(config-ospf)#default-cost ethernet 101
Switch:1(config-ospf)#default-cost fast-ethernet 110
Switch:1(config-ospf)#default-cost gig-ethernet 20
Switch:1(config-ospf)#default-cost ten-gig-ethernet 20
Switch:1(config-ospf)#default-cost twentyfive-gig-ethernet 20
Switch:1(config-ospf)#default-cost Forty-gig-ethernet 20
Switch:1(config-ospf)#default-cost hundred-gig-ethernet 20
Switch:1(config-ospf)#default-cost vlan 2
Switch:1(config-ospf)#timers basic holdown 10
Switch:1(config-ospf)#trap enable
Switch:1(config-ospf)#exit
Switch:1(config)#router ospf enable
Switch:1(config)#show ip ospf
```

```
=====
                        OSPF General - GlobalRouter
=====
```

```

RouterId: 192.0.2.2
AdminStat: disabled
VersionNumber: 2
AreaBdrRtrStatus: false
ASBdrRtrStatus: true
Bad-Lsa-Ignore: false
ExternLsaCount: 0
ExternLsaChecksumSum: 0 (0x0)
TOSSupport: 0
OriginateNewLsas: 0
```

```

RxNewLsas: 0
TrapEnable: false
AutoVirtLinkEnable: false
SpfHoldDownTime: 10
Rfc1583Compatibility: disable
Helper mode: enabled

default-metric :
    ethernet - 101
    fast-ethernet - 110
    gig-ethernet - 20
    ten-gig-ethernet - 20
    twentyfive-gig-ethernet - 20
    forty-gig-ethernet - 20
    hundred-gig-ethernet - 20
Vlan - 1

```

Variable definitions

Use the data in the following table to use the **router-id** command.

Variable	Value
<A.B.C.D>	Configures the OSPF router ID IP address, where A.B.C.D is the IP address.

Use the data in the following table to use the **default-cost** command.

Variable	Value
ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. ethernet is for 10 Mb/s Ethernet (default is 100).
fast-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. fast-ethernet is for 100 Mb/s (Fast) Ethernet (default is 10).
forty-gig-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. forty-gig-ethernet is for 40 Gigabit Ethernet (default is 1).
gig-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. gig-ethernet is for Gigabit Ethernet (default is 1).
hundred-gig-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. hundred-gig-ethernet is for 100 Gigabit Ethernet (default is 1).
ten-gig-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. ten-gig-ethernet is for 10 Gigabit Ethernet (default is 1).
twentyfive-gig-ethernet <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.

Table continues...

Variable	Value
	On a channelized 100 Gbps port, the default-cost for each 25 Gbps channel is 1.
vlan <1-65535>	Configures the OSPF default metrics. vlan is for Vlan interfaces (default is 10).

Use the data in the following table to use the **timers basic holddown** command.

Variable	Value
<3-60>	Configures the OSPF hold-down timer value in seconds. The default is 10.

Use the data in the following table to use the **show ip ospf** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring OSPF for a Port or VLAN

Configure OSPF parameters on a port or VLAN so you can control OSPF behavior on the port or VLAN.

Before you begin

- Enable OSPF globally.
- Ensure IP interfaces exist and are enabled.

About this task

To configure OSPF on a VRF instance for a port or VLAN, you configure OSPF on the port or VLAN, and then associate the port or VLAN with the VRF.

Important:

When you enable OSPF on a VLAN or a port, the switch automatically creates an area 0.0.0.0, and advertises it on the specific VLAN or port, by default. To avoid this behavior, you must manually configure the VLAN or port into a properly configured area on the switch.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]} or interface vlan <1-4059>
```

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the OSPF interface area ID:

```
ip ospf area {A.B.C.D}
```

3. Enable OSPF routing:

```
ip ospf enable
```

4. Choose the OSPF update authentication method:

```
ip ospf authentication-type <message-digest|none|sha-1|sha-2|simple>
```

Both sides of an OSPF connection must use the same authentication type and key.

5. If you choose simple, you must configure the password.

```
ip ospf authentication-key WORD<0-8>
```

6. If you choose an authentication key other than simple such as MD5, Sha-1 or Sha-2, you must configure the digest key first and then assign it to the authentication type.

- a. Create the digest-key:

```
ip ospf digest-key <1-255> key WORD<0-16>
```

- b. Assign the newly created digest key to the authentication type:

```
ip ospf authentication-type <message-digest|none|sha-1|sha-2|simple> primary-digest-key <1-255>
```

7. Specify the interface type:

```
ip ospf network <broadcast|nbma|passive>
```

8. Configure the remaining parameters as required, or accept their default values. View the following variable definitions table for more information.

Example

Configure the OSPF interface area ID to 192.0.2.2, enable OSPF routing, choose the OSPF update authentication method as message-digest, and specify the interface type as broadcast.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#ip ospf area 192.0.2.2
Switch:1(config-if)#ip ospf enable
Switch:1(config-if)#ip ospf authentication-type message-digest
Switch:1(config-if)#ip ospf network broadcast
```

Variable Definitions

Use the data in the following table to use the **ip ospf** commands.

Variable	Value
advertise-when-down enable	Enables or disables AdvertiseWhenDown. If enabled, OSPF advertises the network on this interface as up, even if the port is down. The default is disabled. After you configure a port with no link and enable advertise-when-down, OSPF does not advertise the route until the port is active. OSPF advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
area {A.B.C.D}	Configures the OSPF identification number for the area, typically formatted as an IP address.
authentication-key WORD<0-8>	Configures the eight-character simple password authentication key for the port or VLAN.
authentication-type <message-digest none sha-1 sha-2 simple>	Specifies the type of authentication required for the interface. • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
cost <0-65535>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.
dead-interval <0-2147483647>	Configures the router OSPF dead interval—the number of seconds the OSPF neighbors of a switch must wait before they assume the OSPF router is down. The default is 40. The value must be at least four times the hello interval.
enable	Enables OSPF on the port or VLAN.
hello-interval <1-65535>	Configures the OSPF hello interval, which is the number of seconds between hello packets sent on this interface. The default is 10.
message-digest-key <1-255> md5 WORD<0-16>	Configures the MD5 key, you can configure a maximum of two MD5 keys for an interface. If you configure two keys, the interface uses only the first key. To transition to the second key, configure a primary-md5-key to use the ID of the second configured key, and then delete the first key.

Table continues...

Variable	Value
	! Important: Use the correct key id when two keys are configured. The key id and md5 password must match with the other OSPF routers, to form the OSPF adjacencies. <1-255> is the ID for the MD5 key WORD<0-16> is an alphanumeric password of up to 16 bytes {string length 0–16}
primary-digest-key <1-255>	Use this parameter to transition to a new MD5 key; It changes the primary key used to encrypt outgoing packets. <1-255> is the ID for the new MD5 key.
mtu-ignore enable	Enables maximum transmission unit (MTU) ignore. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
network <broadcast nbma passive>	Specifies the type of OSPF interface.
poll-interval <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
priority <0-255>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you configure the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
retransmit-interval <0-3600>	Configures the retransmit interval for the virtual interface, the number of seconds between link-state advertisement retransmissions.
transit-delay <0-3600>	Configures the transit delay for the virtual interface, which is the estimated number of seconds required to transmit a link-state update over the interface.
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This variable applies only to VLAN interfaces, not ports.

Viewing OSPF errors on a port

Check OSPF errors for administrative and troubleshooting purposes.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display extended information about OSPF errors for the specified port or for all ports:

```
show ip ospf port-error [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the **show ip ospf port-error** command.

Variable	Value
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}</i>	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vrf WORD<1-16>	Specifies the VRF by name.
vrfids WORD<0-512>	Specifies a range of VRFs by ID number.

Configuring OSPF areas on the router

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA).

Before you begin

- Ensure that the VLAN exists if you configure OSPF on a VLAN.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About this task

Place stubby or NSSAs at the edge of an OSPF routing domain. Ensure that you configure all routers in a stubby or NSSA as stubby or NSSA, respectively.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create an OSPF area:

```
area {A.B.C.D}
```

3. Specify the area type:

```
area {A.B.C.D} import <external|noexternal|nssa>
```

4. Configure other OSPF area parameters as required.
5. Ensure that the configuration is correct:

```
show ip ospf area [vrf WORD<1-16>] [vrfrids WORD<0-255>]
```

Example

Create the OSPF area 192.0.2.10, and specify the area type as NSSA. Configure the area support to import summary advertisements into a stub area and configure the import external option for this area as stub. Ensure the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#area 192.0.2.10
Switch:1(config-ospf)#area 192.0.2.10 import nssa
Switch:1(config-ospf)#area 192.0.2.10 import stub
Switch:1(config-ospf)#area 192.0.2.10 import-summaries enable
Switch:1(config-ospf)#show ip ospf area
```

```
=====
                        OSPF Area - GlobalRouter
=====
AREA_ID          STUB_AREA  NSSA          IMPORT_SUM  ACTIVE_IFCNT
-----
192.0.2.10       true         false         true         2

STUB_COST  INTRA_AREA_SPF_RUNS  BDR_RTR_CNT  ASBDR_RTR_CNT  LSA_CNT  LSACK_SUM
-----
0           8                0             0              6        126180
```

Variable Definitions

Use the data in the following table to use the **area {A.B.C.D}** command.

Variable	Value
default-cost <0-16777215>	Specifies the stub area default metric for this stub area, which is the cost from 0–16777215. This metric value applies at the indicated type of service.
import <external noexternal nssa>	Specifies the type of area: - external—stub and NSSA are both false - noexternal—configures the area as stub area. - nssa—configures the area as NSSA.
import-summaries enable	Configures the area support to import summary advertisements into a stub area. Use this variable only if the area is a stub area.
stub	Configures the import external option for this area as stub. A stub area has only one exit point (router interface) from the area.

Use the data in the following table to use the **show ip ospf area** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF.
vrfrids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF area information

View the OSPF area information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF area information:

```
show ip ospf area [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

Example

View the OSPF area information:

```
Switch:1>enable
Switch:1#show ip ospf area
```

```
=====
                        OSPF Area - GlobalRouter
=====
AREA_ID          STUB_AREA  NSSA          IMPORT_SUM  ACTIVE_IFCNT
-----
192.0.2.11       false      false         true         2

STUB_COST  INTRA_AREA_SPF_RUNS  BDR_RTR_CNT  ASBDR_RTR_CNT  LSA_CNT  LSACK_SUM
-----
0           9                0             0             6        117671
```

Variable definitions

Use the data in the following table to use the **show ip ospf area** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfrids WORD<0-512>	Specifies a range of VRF IDs.

Configuring OSPF aggregate area ranges on the router

Use aggregate area ranges to reduce the number of link-state advertisements required within the area. You can also control advertisements.

Before you begin

- Enable OSPF globally.
- Ensure that an area exists.
- You configure OSPF area ranges on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Configure an OSPF area range:

```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
```

3. Configure the advertised metric cost:

```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
advertise-metric <0-65535>
```

4. Configure the advertisement mode:

```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
advertise-mode <summarize|suppress|no-summarize>
```

5. Ensure that the configuration is correct:

```
show ip ospf area-range [vrf WORD<1-16>] [vrfsids WORD<0-512>]
```

Example

Configure an OSPF area range to 192.0.2.2, configure the advertised metric cost to 10, and the advertisement mode to summarize.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link advertise-
metric 10
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link advertise-mode
summarize
```

Variable definitions

Use the data in the following table to use the **area range** command.

Variable	Value
{A.B.C.D} {A.B.C.D/X}	{A.B.C.D} identifies an OSPF area and {A.B.C.D/X} is the IP address and subnet mask of the range, respectively.
advertise-metric <0-65535>	Changes the advertised metric cost of the OSPF area range.
advertise-mode <summarize suppress no-summarize>	Changes the advertisement mode of the range.
<summary-link nssa-extlink>	Specifies the link-state advertisement (LSA) type. If you configure the range as type nssa-extlink, you cannot configure the advertise-metric.

Use the data in the following table to help you use the **show ip ospf area-range** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF area range information

View the OSPF area range information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF area range information:

```
show ip ospf area-range [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the **ip ospf area-range** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Enabling automatic virtual links

Use automatic virtual links to provide an automatic, dynamic backup link for vital OSPF traffic. Automatic virtual links require more system resources than manually configured virtual links.

Before you begin

- You configure automatic virtual links on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Enable the automatic virtual links feature for the router:

```
auto-vlink
```

Configuring an OSPF area virtual interface

Use manual virtual interfaces to provide a backup link for vital OSPF traffic with a minimum of resource use.

Before you begin

- Enable OSPF globally.
- You configure an OSPF area virtual interface on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf`. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About this task

Both sides of the OSPF connection must use the same authentication type and key.

You cannot configure a virtual link using a stub area or an NSSA.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create an OSPF area virtual interface:

```
area virtual-link {A.B.C.D} {A.B.C.D}
```

3. Choose the OSPF update authentication method:

```
area virtual-link {A.B.C.D} {A.B.C.D} authentication-type <message-
digest|none|sha-1|sha-2|simple>
```

Both sides of an OSPF connection must use the same authentication type and key.

4. If required, configure an MD5 key for the virtual interface:

```
area virtual-link message-digest-key {A.B.C.D} {A.B.C.D} <1-255>
md5-key WORD<1-16>
```

5. Configure optional parameters, as required.

6. Ensure that the configuration is correct:

```
show ip ospf virtual-link {A.B.C.D} {A.B.C.D} [vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

Example

Create an OSPF area virtual interface with an area ID of 192.0.2.12 and the virtual interface ID of 198.51.100.2, choose the OSPF update authentication method to simple, and the hello-interval to 100.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2
authentication-type simple
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2 hello-
interval 100
```

Variable Definitions

Use the data in the following table to use the **area virtual-link** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the area ID and the virtual interface ID.
authentication-key WORD<0-8>	Configures the authentication key of up to eight characters.
authentication-type <message-digest none sha-1 sha-2 simple>	Specifies the type of authentication required for the interface. • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
dead-interval <0-2147483647>	Configures the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
hello-interval <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
primary-digest-key <1-255>	Use this parameter to transition to a new MD5 key; It changes the primary key used to encrypt outgoing packets.

Table continues...

Variable	Value
	<1-255> is the ID for the MD5 key.
retransmit-interval <0-3600>	Configures the retransmit interval for the virtual interface, the number of seconds between LSA retransmissions. The range is from 1–3600.
transit-delay <0-3600>	Configures the transit delay for the virtual interface, the estimated number of seconds required to transmit a link-state update over the interface. The range is from 1–3600.

Use the data in the following table to use the **area virtual-link message-digest-key** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the area ID and the virtual interface ID.
<1-255>	Specifies the ID for the message digest key
md5-key WORD<1–16>	Configures the MD5 key, you can configure a maximum of two MD5 keys for an interface. If you configure two keys, the interface uses only the first key. To transition to the second key, configure a primary-md5-key to use the ID of the second configured key, and then delete the first key.  Important: Use the correct key id when two keys are configured. The key id and md5 password must match with the other OSPF routers, to form the OSPF adjacencies. WORD<1–16> is an alphanumeric password of up to 16 characters.

Use the data in the following table to use the **show ip ospf virtual-link** command.

Variable	Value
<A.B.C.D> <A.B.C.D>	Specifies the area ID and the virtual interface ID.
vrf WORD<1-16>	Specifies a VRF.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring an OSPF area on a VLAN or port

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or NSSA. Place stubby or NSSAs at the edge of an OSPF routing domain.

Before you begin

- Enable OSPF globally.
- Ensure that the VLAN exists.

About this task

Ensure that you configure all routers in a stubby or NSSA as stubby or NSSA, respectively.

To configure OSPF areas on a VRF instance for a port or VLAN, you configure OSPF on the port or VLAN, and then associate the port or VLAN with the VRF.

Procedure

1. Enter Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-
port]][, ...]} or interface vlan <1-4059>
```

* Note:

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create an OSPF area on the VLAN or port:

```
ip ospf area {A.B.C.D}
```

3. Specify the type of network:

```
ip ospf network <broadcast|nbma|passive>
```

4. Configure other OSPF area parameters as required.

Example

Create an OSPF area 192.0.2.2 on VLAN 1, and specify the type of network as broadcast.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#ip ospf area 192.0.2.2
Switch:1(config-if)#ip ospf network broadcast
```

Variable Definitions

Use the data in the following table to help you use the **ip ospf** command.

Variable	Value
{A.B.C.D}	Specifies the area ID.
authentication-key WORD<0-8>	Configures the eight-character simple password authentication key for the port or VLAN.
authentication-type <message-digest none sha-1 sha-2 simple>	Specifies the type of authentication required for the interface. • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter.

Table continues...

Variable	Value
	- MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. - sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. - sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
cost <0-65535>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.
dead-interval <0-2147483647>	Configures the the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
hello-interval <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
mtu-ignore enable	Enables MTU ignore. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
network <broadcast nbma passive>	Specifies the type of OSPF interface.
poll-interval <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
primary-digest-key <1-255>	Use this parameter to transition to a new MD5 key; it changes the primary key used to encrypt outgoing packets. <1-255> is the ID for the message digest key.
priority <0-255>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you set the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
retransmit-interval <0-3600>	Configures the retransmit interval: the number of seconds between LSA retransmissions. The range is from 1–3600.
transit-delay <0-3600>	Configures the transit delay: the estimated number of seconds it takes to transmit a link-state update over the interface. The range is from 1–3600.

Configuring an OSPF host route

Configure host routes when the switch resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

Before you begin

- Globally enable OSPF.
- You configure an OSPF host route on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf`. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About this task

Use a host route to create a custom route to a specific host to control network traffic.

You can specify which hosts directly attach to the router, and the metrics and types of service to advertise for the hosts.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create a host route:

```
host-route {A.B.C.D} [metric <0-65535>]
```

3. Ensure that the configuration is correct:

```
show ip ospf host-route [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Create a host route on IP address 192.0.2.20 with a metric of 20.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#host-route 192.0.2.20 metric 20
Switch:1(config-ospf)#show ip ospf host-route
Switch:1(config-ospf)#show ip ospf host-route
```

```
=====
                        OSPF Host Route - GlobalRouter
=====
HOSTIPADDR      TOS  METRIC
-----
192.0.2.20      -    20
```

Variable definitions

Use the data in the following table to use the **host-route** command.

Variable	Value
{A.B.C.D}	Specifies the IP address of the host router in a.b.c.d format.
metric <0-65535>	Configures the metric (cost) for the host route.

Use the data in the following table to use the **show ip ospf host-route** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring OSPF NBMA neighbors

Configure NBMA neighbors so that the interface can participate in designated router election. All OSPF neighbors that you manually configure are NBMA neighbors.

Before you begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Ensure that the interface is NBMA.
- You configure OSPF NBMA neighbors on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create an NBMA OSPF neighbor:

```
neighbor {A.B.C.D} priority <0-255>
```

3. Ensure that the configuration is correct:

```
show ip ospf neighbor [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Create an NBMA OSPF neighbor.

```
Switch:1>enable
Switch:1#configure terminal
```

```
Switch:1(config)#router ospf
Switch:1(config-ospf)#neighbor 198.51.100.2 priority 10
```

Variable definitions

Use the data in the following table to use the **neighbor** command.

Variable	Value
{A.B.C.D}	Identifies an OSPF area in IP address format a.b.c.d.
priority <0-255>	Changes the priority level of the neighbor.

Use the data in the following table to use the **show ip ospf neighbors** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Enabling or disabling Helper mode for OSPFv2

About this task

By default, OSPF Helper mode is enabled when OSPF is configured. You can disable helper mode by the following command and re-enable it again by using no or default operators.

Procedure

1. Enter OSPF Router Configuration mode:


```
enable
configure terminal
router ospf
```
2. Enter the following command to disable Helper mode:


```
helper-mode-disable
```
3. Enter the following command to enable Helper mode:


```
no helper-mode-disable
```

Or

```
default helper-mode disable
```

Example

Disable Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#helper-mode-disable
```

Enable Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#no helper-mode-disable
```

Applying OSPF route acceptance policies

Use a route policy to define how the switch redistributes external routes from a specified source into an OSPF domain. The policy defines which route types the switch accepts and redistributes.

Before you begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that the area exists.
- You apply OSPF route acceptance policies on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf`. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create an acceptance policy instance:

```
accept adv-rtr {A.B.C.D}
```

3. Configure the type of metric to accept:

```
accept adv-rtr {A.B.C.D} metric-type <type1|type2|any>
```

4. Indicate the route policy:

```
accept adv-rtr {A.B.C.D} route-map WORD<0-64>
```

5. Enable a configured OSPF route acceptance instance:

```
accept adv-rtr {A.B.C.D} enable
```

6. Ensure that the configuration is correct:

```
show ip ospf accept [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Create an acceptance policy instance, configure the type of metric to accept, indicate the route policy and enable the OSPF route acceptance instance. Ensure the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#router ospf
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 metric-type type1
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 route-map test1
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 enable
Switch:1(config-ospf)#show ip ospf accept
=====
                        Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE  ENABLE  POLICY
-----
192.0.2.11   type1      true    test1
```

Variable definitions

Use the data in the following table to use the **accept adv-rtr** command.

Variable	Value
<A.B.C.D>	Specifies the IP address.
enable	Enables an OSPF acceptance policy.
metric-type <type1 type2 any>	Configures the metric type as type 1, type 2, or any.
route-map WORD<0-64>	Configures the route policy by name.

Use the data in the following table to use the **ip ospf accept** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF configuration information

View the OSPF configuration information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF configuration information:

```
show ip ospf accept [vrf WORD<1-16>] [vrffids WORD<0-512>]
```

Example

```
Switch:1#show ip ospf accept
```

```
=====
                        Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE  ENABLE  POLICY
-----
192.0.2.11   type1     true   test1
```

Variable definitions

Use the data in the following table to use the **show ip ospf accept** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrffids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF link-state database

View the area advertisements and other information in the LSDB to ensure correct OSPF operations.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF link-state database:

```
show ip ospf lsdb [adv_rtr {A.B.C.D}] [area {A.B.C.D}>] [lsa-type
<0-7>] [lsid {A.B.C.D}] [vrf WORD<1-16>] [vrffids WORD<0-512>]
[detail]
```

Example

```
Switch(config-ospf)#show ip ospf lsdb
```

```
=====
                        OSPF LSDB - GlobalRouter
=====

Router Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
```

```

Router      192.0.2.0      192.0.2.0      617  0x80000031  0xeafd
Router      198.51.100.0   198.51.100.0   1033 0x80000030  0xa5f2

          Network Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID    ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
Network     100.1.1.2          192.0.2.0        617  0x8000002f  0xd038

          Summary Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID    ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----

          AsSummary Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID    ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----

          NSSA Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID    ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----

=====
                          AsExternal Lsas
=====

LSTYPE      LINKSTATEID    ADV_ROUTER      ETYPE  METRIC  ASE_FWD_ADDR      AGE  SEQ_NBR
CSUM
-----

```

Variable definitions

Use the data in the following table to use the **show ip ospf lsdb** command.

Variable	Value
adv_rtr {A.B.C.D}	Specifies the advertising router.
area {A.B.C.D}	Specifies the OSPF area.
detail	Provides detailed output.
lsa-type <0-7>	Specifies the link-state advertisement type in the range of 0–7.
lsid {A.B.C.D}	Specifies the link-state ID.
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF external link-state database

View the LSDB to determine externally learned routing information. Information appears for all metric types or for the type you specify.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View the OSPF autonomous system external (ASE) link-state advertisements:
`show ip ospf ase [metric-type <1-2>] [vrf WORD<1-16>] [vrfrids WORD<0-512>]`

Example

```
Switch:1#show ip ospf ase
=====
                        OSPF AsExternal Lsas - GlobalRouter
=====
LSTYPE      LINKSTATEID      ADV_ROUTER      ETYPE  METRIC  ASE_FWD_ADDR      AGE  SEQ_NBR
CSUM
=====
```

Variable definitions

Use the data in the following table to use the `show ip ospf ase` command.

Variable	Value
metric-type <1-2>	Specifies the metric type.
vrf WORD<1-16>	Identifies the VRF by name.
vrfrids WORD<0-512>	Specifies a VRF by ID.

Configuring route redistribution to OSPF

Configure a redistribute entry to announce certain routes into the OSPF domain, including DvR host routes, static routes, direct routes, Routing Information Protocol (RIP) routes, OSPF routes, IS-IS routes or Border Gateway Protocol (BGP) routes. Optionally, use a route policy to control the redistribution of routes.

Before you begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that you set OSPF as the boundary router.

*** Note:**

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create the redistribution instance:

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> [vrf-src
WORD<0-16>]
```

3. Apply a route policy if required:

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> route-map
WORD<0-64> [vrf-src WORD<0-16>]
```

4. Configure other parameters, as required.

5. Enable the redistribution.

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> enable [vrf-src
WORD<0-16>]
```

6. Ensure that the configuration is correct:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrfs WORD<0-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution.

```
ip ospf apply redistribute <bgp|direct|isis|ospf|rip|static|dvr>
[vrf WORD<1-16>] [vrf-src WORD<0-16>]
```

Changes do not take effect until you apply them.

9. View all routes (including DvR host routes) that are redistributed into OSPF:

- a. View the routes that are redistributed from the GRT to OSPF:

```
show ip ospf lsdbs
```

- b. View the routes that are redistributed to OSPF for a specific VRF instance:

```
show ip ospf lsdbs [vrf WORD<1-64>] [vrfs WORD<0-512>]
```

Example

Example 1:

Redistribute static routes from the GRT to OSPF.

Create the redistribution instance, apply a route policy, enable redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#redistribute static

WARNING: Routes will not be injected until apply command is issued after enable command
Switch:1(config-ospf)#redistribute static route-map policy1
Switch:1(config-ospf)#redistribute static enable
Switch:1(config-ospf)#exit
Switch:1(config)#ip ospf apply redistribute static
Switch:1(config)#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====
```

SRC-VRF	SRC	MET	MTYPE	SUBNET	ENABLE	RPOLICY
GlobalRouter	STAT	0	type2	allow	TRUE	policy1
GlobalRouter	DVR	0	type2	allow	FALSE	

Example 2:

Redistribute DvR host routes from the GRT to OSPF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#redistribute dvr
Switch:1(config-ospf)#redistribute dvr enable
Switch:1(config-ospf)#exit
Switch:1(config)#ip ospf apply redistribute dvr
Switch:1(config)#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====
```

SRC-VRF	SRC	MET	MTYPE	SUBNET	ENABLE	RPOLICY
GlobalRouter	DVR	0	type2	allow	TRUE	test1

Example 3:

Redistribute DvR host routes to OSPF for a specific VRF vrf1:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrf1
```

```
Switch:1(router-vrf)#ip ospf redistribute dvr
Switch:1(router-vrf)#ip ospf redistribute dvr enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip ospf apply redistribute dvr vrf vrf1
```

View the DvR host routes that are distributed into OSPF:

```
Switch:1(config)#show ip ospf lsdb vrf vrf1
```

```
=====
                        OSPF LSDB - VRF vrf1
=====

Router Lsas in Area 0.0.0.0
LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----
Router    192.0.2.1            192.0.2.1      1603 0x8000002d 0xf226
Router    203.0.113.1          203.0.113.1    1608 0x8000002a 0x4104

Network Lsas in Area 0.0.0.0
LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----
Network    14.1.1.11        192.0.2.1      1635 0x80000003 0x4909

Summary Lsas in Area 0.0.0.0
LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----

AsSummary Lsas in Area 0.0.0.0

LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----

NSSA Lsas in Area 0.0.0.0
LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----

Opaque-Loc Lsas in Area 0.0.0.0
LSTYPE    LINKSTATEID    ADV_ROUTER    AGE  SEQ_NBR    CSUM
-----

=====
                        AsExternal Lsas
=====
LSTYPE    LINKSTATEID    ADV_ROUTER    ETYPE  METRIC  ASE_FWD_ADDR    AGE  SEQ_NBR    CSUM
-----
AsExternal 101.1.1.3      203.0.113.1    2       1       0.0.0.0        1563 0x80000003 0xe7a7
AsExternal 101.1.1.4      203.0.113.1    2       1       0.0.0.0        1477 0x80000003 0xddb0
AsExternal 102.1.1.3      203.0.113.1    2       1       0.0.0.0        1528 0x80000003 0xdab3
AsExternal 102.1.1.4      203.0.113.1    2       1       0.0.0.0        1480 0x80000003 0xd0bc
AsExternal 103.1.1.3      203.0.113.1    2       1       0.0.0.0        1531 0x80000003 0xcdbf
...
...
...
```

Variable definitions

Use the data in the following table to use the **redistribute** command.

Variable	Value
enable	Enables the OSPF route redistribution instance.

Table continues...

Variable	Value
metric <0–65535>	Configures the metric to apply to redistributed routes.
metric-type <type1 type2>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
route-map WORD<0-64>	Configures the route policy to apply to redistributed routes.
subnets <allow suppress>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
vrf-src WORD<0-16>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
<bgp direct isis ospf rip static dvr>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, dvr or static.

Use the data in the following table to use the **ip ospf apply redistribute** command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance.
vrf-src WORD<0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
WORD<0–32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Viewing the OSPF redistribution configuration information

Displays the OSPF redistribution configuration information.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF redistribution configuration information:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====
SRC-VRF          SRC  MET MTYPE SUBNET  ENABLE RPOLICY
-----
GlobalRouter     STAT 0   type2 allow   TRUE
```

Variable definitions

Use the data in the following table to use the **show ip ospf redistribute** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring interVRF route redistribution for OSPF

Use route redistribution so that a VRF interface can announce routes learned by other protocols, for example, OSPF or BGP. The switch supports interVRF route redistribution. Use a route policy to control the redistribution of routes.

You can also redistribute inter-VRF DvR routes to OSPF.

Before you begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create the redistribution instance:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr>
```

3. Apply a route policy if required:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr> route-map
WORD<0-64> [vrf-src WORD<0-16>]
```

4. Configure other parameters, as required.

5. Enable the redistribution:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr> enable
[vrf-src WORD<0-16>]
```

6. Ensure that the configuration is correct:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrfsids WORD<0-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution:

```
ip ospf apply redistribute <bgp|direct|isis|ospf|rip|static|dvr>
[vrf WORD<1-16>] [vrf-src WORD<0-16>]
```

Example

Example 1:

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf red
Switch:1(router-vrf)#ip ospf redistribute isis
Switch:1(router-vrf)#ip ospf redistribute isis route-map test2
Switch:1(router-vrf)#ip ospf redistribute isis enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip ospf apply redistribute isis
```

Example 2:

This example demonstrates redistribution of inter-VRF routes (both direct and DvR routes) to OSPF, with a route policy configured.

Redistribute inter-VRF DvR routes between VRFs (with VRF IDs 10 and 30), to OSPF.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf 10
Switch:1(router-vrf)#ip prefix-list "test10" 192.0.2.0/24 ge 25 le 32
Switch:1(router-vrf)#route-map "test10" 1
Switch:1(router-vrf)#permit
Switch:1(router-vrf)#enable
Switch:1(router-vrf)#match network "test10"
Switch:1(router-vrf)#set metric 99
Switch:1(router-vrf)#exit

Switch:1(config)#router vrf 30
Switch:1(router-vrf)#ip ospf redistribute direct vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute direct enable vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr route-map "test10" vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr enable vrf-src 10
Switch:1(router-vrf)#exit

Switch:1(config)#ip ospf apply redistribute direct vrf 30 vrf-src 10
Switch:1(config)#ip ospf apply redistribute dvr vrf 30 vrf-src 10
```

Variable definitions

Use the data in the following table to use the **ip ospf redistribute** command.

Variable	Value
enable	Enables the OSPF route redistribution instance.
metric <0–65535>	Configures the metric to apply to redistributed routes.
metric-type <type1 type2>	Specifies a metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
route-map WORD<0-64>	Configures the route policy to apply to redistributed routes.
subnets <allow suppress>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
vrf-src WORD<0-16>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
<bgp direct isis ospf rip static dvr>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, static or dvr.

Use the data in the following table to use the **ip ospf apply redistribute** command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance.
vrf-src WORD<0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
WORD<0–32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Forcing shortest-path calculation updates

Force the switch to update its shortest-path calculations so that the switch uses the latest OSPF routing information. Manually initiate a shortest path first (SPF) run, or calculation, to immediately update the OSPF LSDB. This action is useful in the following circumstances:

- when you need to immediately restore a deleted OSPF-learned route
- when the routing table entries and the LSDB do not synchronize

Before you begin

- You can perform this procedure in one of the following CLI modes: User EXEC, Privileged EXEC, or Global Configuration.

About this task

This process is computationally intensive. Use this command only if required.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Force the router to update its shortest-path calculations:

```
ip ospf spf-run [vrf WORD<1-16>]
```

Example

Force the router to update its shortest-path calculations:

```
Switch:1>ip ospf spf-run
```

Variable definitions

Use the data in the following table to use the `ip ospf spf-run` command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF instance by name.

Viewing the OSPF default cost information

View the OSPF default cost information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF cost information:

```
show ip ospf default-cost [vrf WORD<1-16>] [vrfs WORD<0-512>]
```

Example

View the OSPF cost information on the switch:

```
Switch:1#show ip ospf default-cost vrf 3
```

```
=====
                        OSPF Default Metric - VRF 3
=====
10MbpsPortDefaultMetric: 100
100MbpsPortDefaultMetric: 10
1000MbpsPortDefaultMetric: 1
10000MbpsPortDefaultMetric: 1
25000MbpsPortDefaultMetric: 1
40000MbpsPortDefaultMetric: 1
100000MbpsPortDefaultMetric: 1
```

Variable definitions

Use the data in the following table to use the `show ip ospf default-cost` command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF interface statistics

Use statistics to help you monitor OSPF performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF interface statistics:

```
show ip ospf ifstats [detail] [mismatch] [vlan <1-4059>] [vrf
WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the OSPF interface statistics:

```
Switch:1#show ip ospf ifstats
```

```
=====
                        OSPF Interface Statistics - GlobalRouter
=====
-----HELLOS-----DBS-----LS REQ--LS UPD---LS ACK---
INTERFACE      RX      TX      RX      TX      RX      TX      RX      TX      RX      TX
-----
192.0.2.3       428     431      0       0       0       0       0       0       0       0
192.0.2.11      1454    493     14      13       4       5      66     54     58      3
```

Variable definitions

Use the data in the following table to use the **show ip ospf ifstats** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
detail	Displays the details of the OSPF.
mismatch	Specifies the number of times the area ID is not matched.
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF timer information

Display OSPF timers information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF timers information:

```
show ip ospf int-timers [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1#show ip ospf int-timers
```

=====						
OSPF Interface Timer - GlobalRouter						
=====						
INTERFACE	AREAID	TRANSIT DELAY	RETRANS INTERVAL	HELLO INTERVAL	DEAD INTERVAL	POLL INTERVAL

192.0.2.1	0.0.0.0	1	5	10	40	120
192.0.2.11	0.0.0.0	1	5	10	40	120
192.0.2.3	0.0.0.0	1	5	10	40	120
=====						
Ospf Virtual Interface Timer						
=====						
AREAID	NBRIPADDR	TRANSIT DELAY	RETRANS INTERVAL	HELLO INTERVAL	DEAD INTERVAL	

Variable definitions

Use the data in the following table to use the **show ip ospf int-timers** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF neighbor information

Displays OSPF neighbor information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF neighbor information:

```
show ip ospf neighbor [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View OSPF neighbor information:

```
Switch:1#show ip ospf neighbor
```

```
=====
                        OSPF Neighbors - GlobalRouter
=====
INTERFACE          NBRROUTERID      NBRIPADDR        PRIO_STATE    RTXQLEN  PERM  TTL
-----
192.0.2.5           192.0.2.1        192.0.2.6         1    Full    0      Dyn    31
198.51.100.7        198.51.100.1     198.51.100.8      128  Restart  0      Dyn    147 H
Total ospf neighbors: 2
H = Helping a Restarting neighbor
```

Variable definitions

Use the data in the following table to use the **show ip ospf neighbor** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF authentication information

Display OSPF authentication information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF authentication information:

```
show ip ospf int-auth [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the OSPF authentication information:

```
Switch:1#show ip ospf int-auth

=====
                        OSPF Interface AuthKey - GlobalRouter
=====
INTERFACE      AUTHTYPE AUTHKEY
-----
192.0.2.1      none
192.0.2.11     none
192.0.2.3      none
```

Variable definitions

Use the data in the following table to use the **show ip ospf int-auth** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF performance statistics

Use statistics to help you monitor OSPF performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF performance statistics:

```
show ip ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the OSPF performance statistics:

```
Switch:1#show ip ospf stats

=====
                        OSPF Statistics - GlobalRouter
=====
      NumBufAlloc: 1138
      NumBufFree: 1138
NumBufAllocFail: 0
NumBufFreeFail: 0
      NumTxPkt: 1144
      NumRxPkt: 2287
      NumTxDropPkt: 0
      NumRxDropPkt: 0
      NumRxBadPkt: 0
```

```

    NumSpfRun: 19
    LastSpfRun: 0 day(s), 00:26:15
    LsdbTblSize: 7
    NumAllocBdDDP: 5
    NumFreeBdDDP: 5
    NumBadLsReq: 0
    NumSeqMismatch: 0
    NumOspfRoutes: 7
    NumOspfAreas: 0
    NumOspfAdjacencies: 3

    NumOspfNbrs: 3

```

Variable definitions

Use the data in the following table to use the **show ip ospf stats** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Viewing the OSPF virtual link information

Displays the OSPF virtual link information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF virtual link information:

```
show ip ospf virtual-link {A.B.C.D} {A.B.C.D} [vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

Example

View the OSPF virtual link information:

```
Switch:1#show ip ospf virtual-link
```

```

=====
                        OSPF Interface AuthKey - GlobalRouter
=====
INTERFACE      AUTHTYPE AUTHKEY
-----
192.0.2.11      none

```

Variable definitions

Use the data in the following table to use the **show ip ospf virtual-link** command.

Variable	Value
{A.B.C.D} {A.B.C.D} vrf WORD<1-16>	Specifies the area ID and the virtual interface ID. The first IP address specifies the area ID and the second specifies the virtual interface ID.
{A.B.C.D} {A.B.C.D} vrfids WORD<0-512>	Displays OSPF configuration for a particular VRF. Specifies a VRF by name.
{A.B.C.D} {A.B.C.D}	Specifies a range of VRF IDs.

Viewing the VRF configurations

Use the following command to view VRF configurations.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the VRF configuration:
show ip ospf vrf WORD<1-16>

Example

View the VRF configuration:

```
Switch:1#show ip ospf vrf virtualrandf1
```

```
=====
                        OSPF General - VRF virtualrandf1
=====

      RouterId: 192.0.2.1
      AdminStat: disabled
      VersionNumber: 2
      AreaBdrRtrStatus: false
      ASBdrRtrStatus: false
      Bad-Lsa-Ignore: false
      ExternLsaCount: 0
      ExternLsaCksumSum: 0(0x0)
      TOSSupport: 0
      OriginateNewLsas: 0
      RxNewLsas: 0
      TrapEnable: false
      AutoVirtLinkEnable: false
      SpfHoldDownTime: 10
      Rfc1583Compatibility: disable

      default-metric :
                    ethernet - 100

--More-- (q = quit)
```

Variable definitions

Use the data in the following table to use the **show ip ospf vrf** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.

Viewing the VRFIDS

Use the following command to view VRFIDS.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the VRF IDS:

```
show ip ospf vrfids WORD<0-512>
```

Example

View the VRF IDs:

```
Switch:1#show ip ospf vrfids 1
```

```
=====
                        OSPF General - VRF virtualrandf1
=====

    RouterId: 192.0.2.1
    AdminStat: disabled
    VersionNumber: 2
    AreaBdrRtrStatus: false
    ASBdrRtrStatus: false
    Bad-Lsa-Ignore: false
    ExternLsaCount: 0
    ExternLsaCksumSum: 0 (0x0)
    TOSSupport: 0
    OriginateNewLsas: 0
    RxNewLsas: 0
    TrapEnable: false
    AutoVirtLinkEnable: false
    SpfHoldDownTime: 10
    Rfc1583Compatibility: disable

    default-metric :
        ethernet - 100

--More-- (q = quit)
```

Variable definitions

Use the data in the following table to use the **show ip ospf vrfid** command.

Variable	Value
vrfids WORD<0-512>	Specifies a range of VRF IDs.

OSPF configuration using EDM

Configure Open Shortest Path First (OSPF) parameters so that the switch can participate in OSPF routing operations. The following section describes procedures that you use while you configure OSPF using Enterprise Device Manager (EDM).

Configuring OSPF globally

Configure OSPF parameters, such as automatic virtual links and OSPF metrics, so you can control OSPF behavior on the system.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.
- Assign an IP address to the switch.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **General** tab.
4. Specify the OSPF router ID.
5. In AdminStart, select **enabled**.
6. **(Optional)** If required, configure the metrics that OSPF uses for different link speeds.
The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
7. **(Optional)** To enable the switch to use OSPF SNMP traps, select the **TrapEnable** check box.
8. **(Optional)** To enable the automatic creation of virtual links, select the **AutoVirtLinkEnable** check box.
9. **(Optional)** Configure the OSPF holddown timer as required.
10. Click **Apply**.

General field descriptions

Use the data in the following table to use the **General** tab.

Name	Description
RouterId	Specifies the OSPF router ID. This variable has the same format as an IP address but distinguishes this router from other routers in the OSPF domain.

Table continues...

Name	Description
AdminStat	Shows the administrative status of OSPF for the router. Enabled denotes that the OSPF process is active on at least one interface; disabled disables it for all interfaces. The default is disabled.
VersionNumber	Specifies the OSPF version.
AreaBdrRtrStatus	Denotes if this router is an area border router (ABR). AreaBdrRtrStatus value must be true to create a virtual router interface.
ASBdrRtrStatus	Specifies ASBR status. If you select the ASBdrRtrStatus check box, the router is an autonomous system boundary router (ASBR).
ExternLsaCount	Shows the number of external (LS type 5) link-state advertisements in the link-state database.
ExternLsaCksumSum	Shows the 32-bit unsigned sum of the link-state checksums of the external link-state advertisements in the link-state database. This sum determines if a change occurred in a router link-state database and compares the link-state databases of two routers.
OriginateNewLsas	Shows the number of new link-state advertisements originated from this router. This number increments each time the router originates a new link-state advertisement (LSA).
RxNewLsas	Shows the number of received link-state advertisements that are new instances. This number does not include new instances of self-originated link-state advertisements.
10MbpsPortDefaultMetric	Indicates the default cost applied to 10 Mbps interfaces (ports). The default is 100.
100MbpsPortDefaultMetric	Indicates the default cost applied to 100 Mbps interfaces (ports). The default is 10.
1000MbpsPortDefaultMetric	Indicates the default cost applied to 1 Gbps interfaces (ports). The default is 1.
10000MbpsPortDefaultMetric	Indicates the default cost applied to 10 Gbps interfaces (ports). The default is 1.
25000MbpsPortDefaultMetric	Indicates the default cost applied to 25 Gbps interfaces (channelized 100 Gbps ports). The default is 1.
40000MbpsPortDefaultMetric	Indicates the default cost applied to 40 Gbps interfaces (ports). The default is 1.
100000MbpsPortDefaultMetric	Indicates the default cost applied to 100 Gbps interfaces (ports). The default is 1.
VlanDefaultMetric	Configures the VLAN interfaces default metric. The default is 10.
TrapEnable	Indicates whether to enable traps for OSPF. The default is false.
AutoVirtLinkEnable	Enables or disables the automatic creation of virtual links. The default is false.
SpfHoldDownTime	Specifies the OSPF holddown timer (3–60 seconds). The default is 10 seconds.

Table continues...

Name	Description
	The holddown timer delays a metric change due to a routing table update by x seconds. If you configure the timer to 0, OSPF accepts a new metric change immediately.
OspfAction	Initiates a new Shortest Path First (SPF) run to update the routing table. The default is none.
Rfc1583Compatability	Controls the preference rules used when the router chooses among multiple autonomous system external (ASE) LSAs which advertise the same destination. If enabled, the preference rule is the same as that specified by RFC 1583. If disabled, the preference rule is as described in RFC 2328, which can prevent routing loops when ASE LSAs for the same destination originate from different areas. The default is disable.
LastSpfRun	Indicates the time since the last SPF calculation made by OSPF.
HelperModeDisable	Disables the helper mode. It is enabled by default.

Enabling OSPF globally

Enable OSPF globally enabled to use the protocol on the router. If you disable OSPF globally, all OSPF actions cease.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **General** tab.
4. For **AdminStat**, select the **enabled** or **disabled** option button, as required.
5. Click **Apply**.

Configuring global default metrics

Configure the metrics that OSPF uses for different link speeds. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the **Configuration > IP** folders.
2. Click **OSPF**.
3. Click the **General** tab.
4. Change the metric for one or all of the following:
 - 10MbpsPortDefaultMetric
 - 100MbpsPortDefaultMetric
 - 1000MbpsPortDefaultMetric
 - 10000MbpsPortDefaultMetric
 - 25000MbpsPortDefaultMetric
 - 40000MbpsPortDefaultMetric
 - 100000MbpsPortDefaultMetric
5. Click **Apply**.

Configuring an OSPF interface

Configure OSPF parameters, such as authentication and priority, so you can control OSPF interface behavior. You can specify the interface as passive, broadcast, or Non-Broadcast Multiple Access (NBMA).

Before you begin

- Enable OSPF globally.
- Ensure that the interface exists (the port or VLAN has an IP address).
- You must know the network OSPF password to use password authentication.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Interfaces** tab.
4. Click **Insert**.
5. Select the IP address for the interface from the IP Address list.
6. To designate a router priority, in the **RtrPriority** box, type a new value.
7. In the **Type** area, select the type of OSPF interface you want to create.
8. Select the authentication type you want in the **AuthType** field.

9. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
10. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.
11. Click **Insert**.
12. On the **Interfaces** tab, click **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the Interfaces tab.

Name	Description
IP Address	Specifies the IP address of the current OSPF interface
AddressLessIf	Designates whether an interface has an IP address: Interfaces with an IP address = 0 Interfaces without IP address = ifIndex
Areald	Specifies the OSPF area name in dotted-decimal format. For VLANs, keeping the default area setting on the interface causes link-state database (LSDB) inconsistencies. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdminStat	Specifies the current administrative status of the OSPF interface (enabled or disabled).
State	Specifies the current state of the OSPF interface. The value can be one of the following: • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
RtrPriority	Specifies the OSPF priority to use during the election process for the designated router. The interface with the highest priority becomes the designated router. The interface with the second-highest priority becomes the backup designated router. If the priority is 0, the interface cannot become the designated router or the backup. The range is 0–255. The default is 1.
DesignatedRouter	Specifies the IP address of the designated router.

Table continues...

Name	Description
BackupDesignatedRouter	Specifies the IP address of the backup designated router.
Type	Specifies the type of OSPF interface (broadcast or NBMA).  Note: To make it passive, first create the interface. After interface creation, click VLAN > VLANs to select the VLAN that is created with the OSPF interface. Click the IP tab and select the IP interface that is created with the OSPF interface. Lastly, click the OSPF tab and select Passive for the IfType.
AuthType	Specifies the type of authentication required for the interface. • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
AuthKey	Specifies the key (up to 8 characters) required when you specify simple password authentication in the AuthType parameter.
HelloInterval	Specifies the length of time, in seconds, between hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
TransitDelay	Specifies the length of time, in seconds, required to transmit an LSA update packet over the interface. The default is 1.
RetransInterval	Specifies the length of time, in seconds, required between LSA retransmissions. The default is 5.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the Hello

Table continues...

Name	Description
	interval. To avoid interpretability issues, the RtrDeadInterval value for the OSPF interface must match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. The default is 120.
Events	Indicates the number of times this OSPF interface has changed state, or an error has occurred.

Changing an OSPF non-passive interface type

Change the interface type to designate the interface as either passive, NBMA, or broadcast.

Before you begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- If the interface is currently an NBMA interface with manually configured neighbors, you must first delete all manually configured neighbors.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Interfaces** tab.
4. To disable the interface, double-click the **AdminStat** cell, and then select **disabled**.
5. Click **Apply**.
6. To change the interface type, double-click the **Type** cell, and then choose the new interface type.
7. Click **Apply**.
8. To enable the interface, double-click the **AdminStat** cell, and then select **enabled**.
9. Click **Apply**.

Changing an OSPF passive interface type

Change the interface type to designate the interface as either passive, NBMA, or broadcast.

Before you begin

- Enable OSPF globally.

- Ensure that the interface uses an IP address.
- If the interface is currently an NBMA interface with manually configured neighbors, you must first delete all manually configured neighbors.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click on the VLAN where the OSPF interface is created.
4. Click **IP**.
5. Select the IP Address where the OSPF interface is created.
6. Click the **OSPF** tab.
7. Clear the **Enable** check box to disable the OSPF interface.
8. Click **Apply**.
9. Modify the interface type to passive.
10. Select the **Enable** check box.
11. Click **Apply**.

Viewing the OSPF advanced interface

View the OSPF advanced interface.

Procedure

1. In the navigation tree, expand the following folders: **Configuration> IP**.
2. Click **OSPF**.
3. Click the **Interface Advanced** tab.

Interface Advanced field description

Use the data in the following table to use the OSPF Interface Advanced tab.

Name	Description
Index	Indicates the Index of the OSPF interface.
Metric	Specifies the metric for the type of service (TOS) on this port. The value of the TOS metric is (10^9 / interface speed). The default is 1. • FFFF—No route exists for this TOS.

Table continues...

Name	Description
	• IPCP links—Defaults to 0. • 0—Use the interface speed as the metric value when the state of the interface is up.
AdvertiseWhenDown	Advertises the network on this port as up, even if the port is down. The default is false. After you configure a port with no link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even if the link is down. To disable advertising based on linkstates, disable AdvertiseWhenDown.
IfMtuIgnore	Specifies whether the interface ignores the global maximum transmission unit (MTU) configuration. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.

Configuring NBMA interface neighbors

Configure NBMA neighbors so that the interface can participate in designated router election. All neighbors that you manually insert on the Neighbors tab are NBMA neighbors.

Before you begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Ensure that the interface type is NBMA.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Neighbors** tab.
4. Click **Insert**.
5. Enter the IP address and priority for the first neighbor.
6. Click **Insert**.
7. Add all required neighbors.
8. Click **Apply**.

Neighbors field descriptions

Use the data in the following table to use the Neighbors tab.

Name	Description
NbrIpAddr	Specifies the neighbor IP address.
AddressLessIndex	Indicates addressed and addressless interfaces. This value is 0 on an interface with an IP address. On addressless interfaces, the corresponding value of ifIndex in the Internet standard management information base (MIB).
NbrRtrId	Specifies the router ID of the neighboring router. The router ID has the same format as an IP address but identifies the router independent of its IP address.
Options	Specifies the bit mask that corresponds to the neighbor options parameter.
Priority	Specifies the priority.
State	Specifies the OSPF interface state.
Events	Specifies the number of state changes or error events that occur between the OSPF router and the neighbor router.
Retransmission Queue Length	Specifies the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
ospfNbmaNbrPermanence	Indicates whether the neighbor is a manually configured NBMA neighbor; permanent indicates it is an NBMA neighbor.
HelloSuppressed	Indicates whether hello packets to a neighbor are suppressed.

Configuring OSPF interface metrics

Configure the metrics associated with the peer layer interface to control OSPF behavior. For finer control over port-specific metric speed, you can specify the metric speed when you configure OSPF on a port.

Before you begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **If Metrics** tab.
4. Double-click the value cell, and type a new value.

5. Click **Apply**.

When you enable a port for OSPF routing, the default metric in the port tab is 0. A value of 0 means that the port uses the default metrics for port types that you specify on the OSPF General tab.

If Metrics field descriptions

Use the data in the following table to use the If Metrics tab.

Name	Description
IP Address	Specifies the IP address of the device used to represent a point of attachment in a TCP/IP internetwork.
AddressLessIf	Indicates addressed and addressless interfaces. This variable is 0 on interfaces with IP addresses and equals ifIndex for interfaces that have no IP address.
TOS	Specifies the type of service (TOS). The TOS is a mapping to the IP type of service flags as defined in the IP forwarding table management information base (MIB).
Value	Indicates the metric from the OSPF router to a network in the range.
Status	Specifies the status of the interface as active or not active. This variable is read-only.

Viewing all OSPF-enabled interfaces

View all OSPF-enabled interfaces to determine which interfaces use OSPF routing.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Interfaces** tab.
4. To ensure the latest information appears, click **Refresh**.

Configuring OSPF on a port

Configure OSPF parameters on a port so you can control OSPF behavior on the port.

Important:

When you enable OSPF on a port, the switch automatically creates an area 0.0.0.0, and advertises it on the specific port, by default. To avoid this behavior, you must manually configure the port into a properly configured area on the switch.

Before you begin

- Enable OSPF globally .
- Ensure that the port uses an IP address.
- Ensure that the ospf_md5key.txt file is on the switch to use MD5 authentication.
- You must know the network OSPF password to use password authentication.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
3. Click **IP**.
4. Click the **OSPF** tab.
5. Select the **Enable** check box.
6. Specify the hello interval.
7. Specify the router dead interval.
8. Designate a router priority.
9. Configure a metric.
10. If you want, select an authentication type.
11. If you select **simplePassword** authentication, type a password in the **AuthKey** box.
12. Configure the area ID.
13. If desired, select the **AdvertiseWhenDown** check box.
14. Select an interface type.
15. Type a value in the **PollInterval** box.
16. In the IfMtuIgnore area, select either **enable** or **disable**.
17. Click **Apply**.

OSPF Field Descriptions

Use the data in the following table to use the OSPF tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified port. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds.

Table continues...

Name	Description
	After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet, and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this port in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for the type of service (TOS) on this port. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. - FFFF—No route exists for this TOS. - IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.
AuthType	Specifies the type of authentication required for the interface. - none—Specifies that no authentication required. - simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. - MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. - sha1—Specifies secure hash algorithm (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. - sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
AuthKey	Specifies the key (up to 8 characters) when you specify simple password authentication in the port AuthType variable.

Table continues...

Name	Description
AreaId	Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	Advertises the network on this port as up, even if the port is down. The default is false. After you configure a port with no link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even if the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
IfType	Specifies the type of OSPF interface (broadcast, NBMA, or passive). Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must have the same poll interval.
IfMtuIgnore	Specifies whether the interface ignores the global maximum transmission unit (MTU) configuration. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.

Configuring OSPF on a VLAN

Configure OSPF parameters on a VLAN to control OSPF behavior on the VLAN.

Important:

When you enable OSPF on a VLAN, the switch automatically creates an area 0.0.0.0, and advertises it on the specific VLAN, by default. To avoid this behavior, you must manually configure the VLAN into a properly configured area on the switch.

Before you begin

- Enable OSPF globally.
- Ensure that the VLAN uses an IP address.
- Ensure that the ospf_md5key.txt file is on the switch to use MD5 authentication.
- Ensure that you know the network OSPF to use password authentication.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **OSPF** tab.

The information on the OSPF tab applies only to a routed port or VLAN, which means the VLAN uses an IP address.

7. To enable OSPF on the VLAN interface, select the **Enable** check box.
8. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.
9. To designate a router priority, in the **DesigRtrPriority** box, type the new value.
10. Select the authentication type in the **AuthType** field.
11. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
12. Select the interface type you want to create.
13. Click **Apply**.

OSPF Field Descriptions

Use the data in the following table to use the OSPF tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified VLAN. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.

Table continues...

Name	Description
DesigRtrPriority	Specifies the priority of this VLAN in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for this TOS on this VLAN. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. - FFFF—No route exists for this TOS. - IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.
AuthType	Specifies the type of authentication required for the interface. - none—Specifies that no authentication required. - simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. - MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. - sha1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. - sha-2—Specifies SHA-2, which offers the hash function SHA-256.  Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.
AuthKey	Specifies the key (up to eight characters) when you specify simple password authentication in the VLAN AuthType variable.
Areald	Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	Advertises the network even if the port is down. If true, OSPF advertises the network on this VLAN as up, even if the port is down. The default is false. After you configure a port without a link and enable AdvertiseWhenDown, it does not advertise the route until the port is

Table continues...

Name	Description
	active. Then, OSPF advertises the route even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.
IfType	Specifies the type of OSPF interface (broadcast, NBMA, or passive). Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must use the same poll interval.
IfMtuIgnore	Specifies whether the VLAN ignores the MTU configuration. To allow the switch to accept OSPF DD packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.

Viewing graphs for OSPF on a VLAN

View graphs for OSPF on a VLAN. The graph formats available are: line chart, area chart, bar chart, and pie chart.

Before you begin

- OSPF must be enabled.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN, and then click **IP**.
5. Click the **OSPF** tab.
6. Click **Graph**.
7. **(Optional)** To refresh the values in the table, click **Clear Counters**.
8. To specify the polling interval, from the **Poll Interval** drop down menu, select a value. The options are:

Choice Option	Choice Description
5s	The polling interval is 5 seconds.
10s	The polling interval is 10 seconds.
30s	The polling interval is 30 seconds.
1m	The polling interval is 1 minute.
5m	The polling interval is 5 minutes.

Choice Option	Choice Description
30m	The polling interval is 30 minutes.
1h	The polling interval is 1 hour.

9. Select one or two values.

To select two values, for example, AbsoluteValue and Cumulative. Select the first value, and then press the **Control** key to select the second value. You cannot select more than two values.

10. From the toolbar, click a chart icon. The options are:

Choice Option	Choice Description
Line Chart	Displays a line chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Area Chart	Displays an area chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Bar Chart	Displays a bar chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Pie Chart	Displays a pie chart for the values you selected against the polling interval.

The Chart Legend uses different colors to identify the values you selected that are plotted on the graph.

11. To switch the horizontal and vertical axes values, on the chart toolbar, click **Horizontal**.
12. To switch views of the log scale from high to low values, or low to high values, on the chart toolbar, click **Log Scale**.
13. To switch to another chart using the same values, on the chart toolbar, click a chart icon.

OSPF graph field descriptions

Use the data in the following table to use the OSPF graph tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the OSPF-Graph tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
VersionMismatches	Displays the number version mismatches received by this interface.

Table continues...

Name	Description
AreaMismatches	Displays the number area mismatches received by this interface.
AuthTypeMismatches	Displays the number AuthType mismatches received by this interface.
AuthFailures	Displays the number Authentication failures.
NetMaskMismatches	Displays the number net mask mismatches received by this interface.
HelloIntervalMismatches	Displays the number hello interval mismatches received by this interface.
DeadIntervalMismatches	Displays the number dead interval mismatches received by this interface.
OptionMismatches	Displays the number options mismatches received by this interface.
RxHellos	Displays the number hello packets received by this interface.
RxDBDescrs	Displays the number database descriptor packets received by this interface.
RxLSUpdates	Displays the number Link state update packets received by this interface.
RxLSReqs	Displays the number Link state request packets received by this interface.
RxLSAcks	Displays the number Link state acknowledge packets received by this interface.
TxHellos	Displays the number hello packets transmitted by this interface.
TxDBDescrs	Displays the number database descriptor packets transmitted by this interface.
TxLSUpdates	Displays the number Link state update packets transmitted by this interface.
TxLSReqs	Displays the number Link state request packets transmitted by this interface.
TxLSAcks	Displays the number Link state acknowledge packets transmitted by this interface.

Creating stubby or not-so-stubby OSPF areas

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA).

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

Place stubby areas or NSSAs at the edge of an OSPF routing domain. Ensure that you configure all routers in the stubby or NSSA as stubby or NSSA, respectively.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Areas** tab.

The backbone ID has an area ID of 0.0.0.0.

4. Click **Insert**.
5. Configure the area ID.
6. Select an option in the ImportAsExtern area.

To add a not-so-stubby (NSSA) area, select **importNssa**. To import external LSAs (create a normal OSPF area), select **importExternal**. To not import external LSAs (create a stubby area), select **importNoExternal**.

7. Click **Apply**.

Areas field descriptions

Use the data in the following table to use the Areas tab.

Name	Description
AreaId	Specifies a 32-bit integer that uniquely identifies an area. Area ID 0.0.0.0 is the OSPF backbone. For VLANs, using the default area on the interface causes LSDB inconsistencies.
ImportAsExtern	Specifies the method to import ASE link-state advertisements. The value can be importExternal (default), importNoExternal, or importNssa.
SpfRuns	Specifies the number of SPF calculations performed by OSPF.
AreaBdrRtrCount	Specifies the number of area border routers reachable within this area. Each SPF pass calculates this value, initially zero.
AsBdrRtrCount	Specifies the number of autonomous system border routers reachable within this area. Each SPF pass calculates this value, initially zero.
AreaLsaCount	Specifies the total number of link state advertisements in this area LSDB, excluding AS-external LSAs.
AreaLsaCksumSum	Specifies the number of link-state advertisements. This sum excludes external (LS type 5) link-state advertisements. The sum determines if a change occurred in a router LSDB and compares the LSDB of two routers.

Table continues...

Name	Description
AreaSummary	Specifies whether to send summary advertisements in a stub area.
ActiveifCount	Specifies the number of active interfaces in this area.

Configuring stub area metrics advertised by an ABR

Configure metrics to control the use of routes in a routing domain.

Before you begin

- Enable OSPF globally.
- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Stub Area Metrics** tab.
4. Double-click the metric value to edit it and specify a new metric speed for the required stub areas.
5. Click **Apply**.

Stub Area Metrics field descriptions

Use the data in the following table to use the Stub Area Metrics tab.

Name	Description
Areald	Specifies the 32-bit identifier for the stub area.
TOS	Specifies the type of service associated with the metric.
Metric	Specifies the metric value applied at the indicated type of service. By default, the value equals the lowest metric value at the type of service among the interfaces to other areas.
Status	Specifies the status of the stub area. This variable is read-only.

Inserting OSPF area aggregate ranges

Use aggregate area ranges to reduce the number of link-state advertisements required within the area. You can also control advertisements.

Before you begin

- Enable OSPF globally.

- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Area Aggregate** tab.
4. Click **Insert**.
5. Type the area ID.
6. Select the type of link-state database.
7. Type the IP address of the network.
8. Type the subnet mask.
9. Select the effect.
10. In the **AdvertiseMetric** box, type a cost to advertise for the OSPF area range.
11. Click **Insert**.

Area Aggregate Field Descriptions

Use the data in the following table to use the Area Aggregate tab.

Name	Description
AreaID	Specifies the area in which the address exists.
LsdbType	Specifies the LSDB type: • summaryLink—aggregated summary link • nssaExternalLink—not so stubby area link
IP Address	Specifies the IP address of the network or subnetwork indicated by the range.
Mask	Specifies the network mask for the area range.
Effect	Specifies advertisement methods: • advertiseMatching means advertise the aggregate summary LSA with the same LSID. • doNotAdvertiseMatching means suppress all networks that fall within the entire range. • advertiseDoNotAggregate means advertise individual networks.
AdvertiseMetric	Changes the advertised metric cost for the OSPF area range.

Enabling automatic virtual links

Use automatic virtual links to provide an automatic, dynamic backup link for vital OSPF traffic.

Before you begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **General** tab.
4. Select the **AutoVirtLinkEnable** check box.
5. Click **Apply**.

Configuring a manual virtual interface

Use manual virtual links (interfaces) to provide a backup link for vital OSPF traffic with a minimum of resource use.

Before you begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Virtual If** tab.
4. Click **Insert**.
5. Specify the area ID of the transit area.

The transit area is the common area between two ABRs.

6. Specify the neighbor ID.

The neighbor ID is the IP router ID of the ABR that the other ABR needs to reach the backbone.

7. Click **Insert**.
8. To verify that the virtual link is active, click **Refresh** and check the **State** column.

If the state is point-to-point, the virtual link is active. If the state is down, the virtual link configuration is incorrect.

Virtual If field descriptions

Use the data in the following table to use the Virtual If tab.

Name	Description
Areaid	Specifies the transit area ID that the virtual link traverses.
Neighbor	Specifies the router ID of the virtual neighbor.
TransitDelay	Specifies the estimated number of seconds required to transmit a link-state update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between link-state advertisement, and retransmissions for adjacencies that belong to this interface. This variable also applies to DD and link-state request packets. This value must exceed the expected round-trip time. The default is 5.
HelloInterval	Specifies the length of time, in seconds, between the hello packets that the router sends on the interface. This value must be the same for the virtual neighbor. The default is 10.
RtrDeadInterval	Specifies the number of seconds that expires before neighbors declare the router down. This value must be a multiple of the hello interval. This value must be the same for the virtual neighbor. The default is 60.
State	Specifies the OSPF virtual interface state.
Events	Specifies the number of state changes or error events on this virtual Link.
AuthType	Specifies the authentication type specified for a virtual interface. You can locally assign additional authentication types. The default is none.
AuthKey	Specifies the authentication password. If AuthType is a simple password, the device adjusts and zeros fill the eight octets. Unauthenticated interfaces need no authentication key, and simple password authentication cannot use a key with more than eight octets.

Viewing virtual neighbors

View virtual neighbors to view the area and virtual link configuration for the neighboring device.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Virtual Neighbors** tab.

Virtual Neighbors field descriptions

Use the data in the following table to use the Virtual Neighbors tab.

Name	Description
Area	Specifies the subnetwork in which the virtual neighbor resides.
RtrId	Specifies the 32-bit integer (represented as an IP address) that uniquely identifies the neighbor router in the autonomous system.
IP Address	Specifies the IP address of the virtual neighboring router.
Options	Specifies the bit mask that corresponds to the neighbor options parameter.
State	Specifies the OSPF interface state.
Events	Specifies the number of state changes or error events that occurred between the OSPF router and the neighbor router.
LsRetransQLen	Specifies the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
HelloSuppressed	Specifies whether hello packets from the neighbor are suppressed.

Configuring host routes

Configure host routes when the switch resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

Before you begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

You can specify which hosts directly connect to the router and the metrics and types of service to advertise for the hosts.

Use a host route to create a custom route to a specific host to control network traffic.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.

3. Click the **Hosts** tab.
4. To insert a new host, click **Insert**.
5. In the **IP Address** box , type the area IP address of the new host.
6. In the **Metric** box, type the metric to advertise.
7. Click **Insert**.
8. Click **Apply**.

Hosts field descriptions

Use the data in the following table to use the Hosts tab.

Name	Description
IpAddress	Specifies the IP address of the host that represents a point of attachment in a TCP/IP internetwork.
TOS	Specifies the type of service of the route.
Metric	Specifies the metric advertised to other areas. The value indicates the distance from the OSPF router to a network in the range.
AreaID	Specifies the area where the host is found. By default, the area that submits the OSPF interface is in 0.0.0.0.

Enabling ASBR status

Enable the ASBR status to make the switch an autonomous system boundary router (ASBR). Use ASBRs to advertise nonOSPF routes into OSPF domains so that the routes pass through the domain. A router can function as an ASBR if one or more of its interfaces connects to a non-OSPF network, for example, Routing Information Protocol (RIP), BGP, or Exterior Gateway Protocol (EGP).

Before you begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

To conserve resources, you can limit the number of ASBRs on your network or specifically control which routers perform as ASBRs to control traffic flow.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **General** tab.
4. Select the **ASBdrRtrStatus** check box.

5. Click **Apply**.

Managing OSPF neighbors

View or delete OSPF neighbors to control OSPF operations.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

The OSPF Hello protocol initiates and maintains neighbor relationships. The exception is that, in an NBMA network, you must manually configure permanent neighbors on each router eligible to become the DR. You can add neighbors for NBMA interfaces, but all other neighbors are dynamically learned.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Neighbors** tab.
4. To delete a manually configured neighbor, select the neighbors with a value of **permanent** in the **ospfNbmaNbrPermanence** column.
5. Click **Delete**.
6. Click **Apply**.

Viewing the link-state database

View the area advertisements and other information in the LSDB to ensure correct OSPF operations.

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Link State Database** tab.

Link State Database field descriptions

Use the data in the following table to use the Link State Database tab.

Name	Description
Areald	Identifies the area. The OSPF backbone uses the area ID 0.0.0.0.
Type	Specifies the OSPF interface type. Broadcast LANs, such as Ethernet and IEEE 802.5, use broadcast; X.25 and similar technologies use NBMA; and links that are point-to-point use pointToPoint.
Lsid	Identifies the piece of the routing domain that the advertisement describes.
RouterId	Identifies the router in the autonomous system.
Sequence	Identifies old and duplicate link-state advertisements.
Age	Specifies the age, in seconds, of the link-state advertisement.
Checksum	Contains the checksum of the complete contents of the advertisement, except for the age parameter. The checksum does not include the age parameter so that advertisement age increments without updating the checksum.

Configuring interVRF route redistribution policies

Configure interVRF route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF, RIP, or BGP. Use a route policy to control the redistribution of routes.

Before you begin

- VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **Policy**.
3. Click the **Route Redistribution** tab.
4. Click **Insert**.
5. Click the ellipsis (...) button near the **DstVrflid** box to select the source and destination VRF IDs.
6. Click the ellipsis (...) button near the **SrcVrflid** box to select the source and destination VRF IDs.
7. In the **Protocol** option box, select the protocol.
8. In the **RouteSource** option box, select the route source.
9. Select **enable**.
10. Click the ellipsis (...) button near the **RoutePolicy** box to choose the route policy to apply to the redistributed routes.

11. Configure other parameters as required.
12. Click **Insert**.
13. Click the **Applying Policy** tab.
14. Select **RedistributeApply**.
15. Click **Apply**.

Route Redistribution field descriptions

Use the data in the following table to use the Route Redistribution tab.

Name	Description
DstVrflid	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrflid	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
Subnets	Indicates that subnets must be advertised individually (applies to OSPF only).

Configuring route redistribution to OSPF

Configure a redistribute entry to announce routes of a certain source protocol type into the OSPF domain, for example, static, RIP, or direct. Optionally, use a route policy to control the redistribution of routes.

Before you begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

Important:

Changing the OSPF redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform this procedure. It is recommended that if you want to change default preferences for an OSPF redistribute context, you must do so before you enable the protocols.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Redistribute** tab.
4. Click **Insert**.
5. Select an option for the route source.
6. Select the **enable** option button.
7. Select a route policy.
8. Configure the metric type.
9. Configure the subnet.
10. Click **Insert**.

Redistribute field descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrflid	Specifies the destination virtual router forwarding instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrflid	Specifies the source VRF instance. You cannot configure this variable.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) to use for detailed redistribution of external routes from a specified source into an OSPF domain. Click the ellipsis (...) button and choose from the list in the dialog box.
Metric	Configures the OSPF route redistribution metric for basic redistribution. The value can be a range from 0–65535. A value of 0 indicates to use the original cost of the route.
MetricType	Configures the OSPF route redistribution metric type. The default is type 2.

Table continues...

Name	Description
	The cost of a type 2 route is the external cost, regardless of the interior cost. A type 1 cost is the sum of both the internal and external costs.
Subnets	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.

Viewing OSPF status

View OSPF status.

Procedure

1. In the navigation pane, expand the following folders: **Configuration**> **IP**.
2. Click **OSPF**.
3. Click the **Stats** tab.

Viewing OSPF status graphs

View OSPF status graphs. The graph formats available are: line chart, area chart, bar chart, and pie chart.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **OSPF**.
3. Click the **Stats** tab.
4. (Optional) To refresh the values in the table, click **Clear Counters**.
5. To specify the polling interval, from the **Poll Interval** drop down menu, select a value. The options are:

Option	Description
5s	The polling interval is 5 seconds.
10s	The polling interval is 10 seconds.
30s	The polling interval is 30 seconds.
1m	The polling interval is 1 minute.
5m	The polling interval is 5 minutes.
30m	The polling interval is 30 minutes.
1h	The polling interval is 1 hour.

6. Select one value; for example, AbsoluteValue or Cumulative.
 - Or, select two values; for example, AbsoluteValue and Cumulative.

To select a second value, press the **Ctrl** key, then select the second value. You cannot select more than two values.

- From the toolbar, click a chart icon. The options are:

Option	Description
Line Chart	Displays a line chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Area Chart	Displays an area chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Bar Chart	Displays a bar chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Pie Chart	Displays a pie chart for the values you selected against the polling interval.

The Chart Legend uses different colors to identify the values you selected that are plotted on the graph.

- To switch the horizontal and vertical axes values, on the chart toolbar, click **Horizontal**.
- To switch views of the log scale from high to low values, or low to high values, on the chart toolbar, click **Log Scale**.
- To switch to another chart using the same values, on the chart toolbar, click a chart icon.

Stats field descriptions

Use the data in the following table to use the OSPF Stats tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
LsdbTblSize	Displays the number of entries in the link state database table.
TxPackets	Displays the number of packets transmitted by OSPF.
RxPackets	Displays the number of packets received by OSPF.
TxDropPackets	Displays the number of packets dropped before transmitted by OSPF.
RxDropPackets	Displays the number of packets dropped before received by OSPF.

Table continues...

Name	Description
RxBadPackets	Displays the number of packets received by OSPF that are bad.
SpfRuns	Displays the total number of SPF calculations performed by OSPF, which includes the number of partial route table calculation for incremental updates.
BuffersAllocated	Displays the number of buffers allocated for OSPF.
BuffersFreed	Displays the number of buffers that are freed by the OSPF.
BufferAllocFailures	Displays the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Displays the number of times that OSPF has failed to free buffers.
Routes	Displays the number of OSPF routes added to RTM.
Adjacencies	Displays how many adjacencies are learned through the interface.
Areas	Displays the number of areas configured.
Nbrs	Indicates the number of OSPF neighbors.
BadLsReqs	Indicates the number of bad link state requests.
SeqMismatches	Indicates the number of sequence mismatched packets.
NumAllocDDP	Indicates the number of database description (DD) packet buffers allocated for OSPF.
NumFreeDDP	Indicates the number of DD packet buffers that are freed by the OSPF.

Forcing shortest-path calculation updates

Manually initiate an SPF run, or calculation, to immediately update the OSPF LSDB. This configuration is useful if

- you need to immediately restore a deleted OSPF-learned route
- the routing table entries and the LSDBs do not synchronize

Before you begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About this task

This process is computationally intensive. Use this command only if required.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Double-click **OSPF**.
3. Click the **General** tab.
4. In the **OspfAction** area, select the **runSpf** option button.
5. Click **Apply**.
6. Click **Yes** to force an SPF run.

After you initiate an SPF run, wait at least 10 seconds before you initiate another SPF run.

Chapter 5: RIP

This chapter provides concepts and configuration procedures for Routing Information Protocol (RIP).

RIP fundamentals

Use the information in these sections to help you understand the Routing Information Protocol (RIP). For more information about the Border Gateway Protocol (BGP), see [Configuring BGP Services for VOSS](#).

Routing Information Protocol

In routed environments, routers communicate with one another to track available routes. Routers can dynamically learn about available routes using the RIP. The switch software implements standard RIP to exchange IP route information with other routers.

RIP uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. Each router advertises routing information by sending a routing information update every 30 seconds (one interval). If RIP does not receive information about a network for 180 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 120 seconds, it removes the network from the routing table.

RIP is a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is the distance from one router to the next. This cost or hop count is the metric (see the following figure).




Figure 9: Hop count or metric in RIP

RIP version 1 (RIPv1) advertises default class addresses without subnet masking. RIP version 2 (RIPv2) advertises class addresses explicitly, based on the subnet mask.

The switch supports RIPv2, which supports variable length subnet masks (VLSM) and triggered router updates. RIPv2 sends mask information. If RIP does not receive information about a network for 90 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 180 seconds (six update intervals), it removes the network from the routing table. You can change the default timers by configuring the RIP interface timeout timer and the holddown timer.

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, the highest metric between two networks can be 15 hops or 15 routers.

RIP and route redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if RIP routes exist in a router and they must travel through a BGP network, configure redistribution of RIP routes through BGP. Redistribution sends RIP routes to a router that uses BGP.

You can redistribute routes

- on an interface basis
- on a global basis between protocols on a single VRF instance (intraVRF)
- between the same or different protocols on different VRF instances (interVRF)

To configure interface-based redistribution, configure a route policy, and then apply it to the interface. Configure the match parameter to the protocol from which to learn routes.

You can redistribute routes on a global basis, rather than for every interface. The switch adds support for global RIP redistribution. Use the `ip rip redistribute` command to accomplish the (intraVRF) redistribution of routes through RIP, so that RIP redistribution occurs globally on all RIP-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

If you configure redistribution globally and on an interface, redistribution through the route policy takes precedence.

You can redistribute routes from a protocol in one VRF to RIP in another VRF. You can use a route policy for redistribution control. If you enable route redistribution between VRF instances, ensure that IP addresses do not overlap.

RIP configuration using CLI

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use the command line interface (CLI) to configure and manage RIP.

Configuring RIP globally

Configure RIP parameters on the switch so you can control RIP behavior on the system.

Before you begin

- You configure RIP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip rip`. The VRF must have an RP Trigger of RIP. Not all parameters are configurable on non0 VRFs.

About this task

All router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

You can configure RIP on interfaces while RIP is globally disabled. This way, you can configure all interfaces before you enable RIP for the switch.

Procedure

1. Enter RIP Router Configuration mode:

```
enable
configure terminal
router rip
```
2. Define the default-import-metric for the switch:

```
default-metric <0-15>
```
3. **(Optional)** Configure one or more timer values:

```
timers basic timeout <15-259200> [holddown <0-360>] [update <1-360>]
```

4. Enable RIP on an IP network:

```
network {A.B.C.D}
```

5. Exit to Global Configuration mode:

```
exit
```

6. After the configuration is complete, enable RIP globally:

```
router rip enable
```

7.

8. Check that your configuration is correct:

```
show ip rip [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Define the default-import-metric as 1, the timeout interval as 180, the holddown time as 120, and the update time as 30. Enable RIP on an IP network, and ensure your configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#router rip
Switch:1(config-rip)#default-metric 1
Switch:1(config-rip)#timers basic timeout 180 holddown 120 update 30
Switch:1(config-rip)#network 192.0.2.11
Switch:1(config-rip)#exit
Switch:1(config)#router rip enable
Switch:1(config)#show ip rip
=====
                        RIP Global - GlobalRouter
=====
Default Import Metric : 1
      HoldDown Time   : 120
        Queries       : 0
          Rip          : Enabled
      Route Changes    : 0
    Timeout Interval  : 180
      Update Time     : 30
```

Variable definitions

Use the data in the following table to use the RIP commands in this procedure.

Variable	Value
default-metric <0-15>	Configures the value of default import metric to import a route into a RIP domain. To announce OSPF internal routes into RIP domain, if the policy does not specify a metric value, the default is used. For OSPF external routes, the external cost is used. The default is 8.
domain <0-39321>	Specifies the RIP domain from 0–39321. The default is 0.
holddown <0-360>	Configures the RIP hold-down timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.

Table continues...

Variable	Value
network {A.B.C.D}	Enables RIP on an IP network.
timeout <15-259200>	Configures the RIP timeout interval. The default is 180.
update <1-360>	Configures the RIP update timer. The update time is the time interval, in seconds, between RIP updates. The default is 30.

Use the data in the following table to use the **show ip rip** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring RIP on an interface

Configure RIP on Ethernet ports and VLANs so that they can participate in RIP routing.

Before you begin

- Assign an IP address to the port or VLAN.
- Configure RIP and enable it globally.
- Configure in and out policies.

About this task

RIP does not operate on a port or VLAN until you enable it both globally and on the port or VLAN.

To configure RIP on a VRF instance for a port or VLAN, you configure RIP on the port or VLAN, and then associate the port or VLAN with the VRF.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} or interface vlan <1-4059>
```

Note:

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Define the cost:

```
ip rip cost <1-15>
```

3. Specify an in policy for filtering inbound RIP packets:

```
ip rip in-policy WORD<0-64>
```

- Specify an out policy for filtering outbound RIP packets:

```
ip rip out-policy WORD<0-64>
```

- Enable RIP:

```
ip rip enable
```

- Specify the send mode:

```
ip rip send version <notsend|rip1|rip1comp|rip2>
```

- Specify the receive mode:

```
ip rip receive version <rip1|rip2|rip1orrip2>
```

- Change other RIP parameters from their default values as required.

Example

The following configuration example shows how to configure the switch (labeled R1) to operate only in RIP version 2 mode.




Figure 10: Configuration example-RIPv2 only

Enable RIPv2 send mode on VLAN 2:

```
Switch:1(config-if)# ip rip send version rip2
```

Enable RIPv2 receive mode on VLAN 2:

```
Switch:1(config-if)# ip rip receive version rip2
```

Repeat these commands on VLAN 3 and the port interfaces.

Variable definitions

Use the data in the following table to use the `ip rip` command.

Variable	Value
advertise-when-down enable	Enables or disables AdvertiseWhenDown. If enabled, RIP advertises the network on this interface as up, even if the port is down. The default is disabled.

Table continues...

Variable	Value
	If you configure a port with no link and enable advertise-when-down, it does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
auto-aggregation enable	Enables or disables automatic route aggregation on the port. If enabled, the switch automatically aggregates routes to their natural mask when an interface in a different class network advertises them. The default is disable.
cost <1-15>	Configures the RIP cost for this port (link).
default-listen enable	Enables DefaultListen. The switch accepts the default route learned through RIP on this interface. The default is disabled.
default-supply enable	Enables DefaultSupply. If enabled, this interface must advertise a default route. The default is false. RIP advertises the default route only if it exists in the routing table.
enable	Enables RIP routing on the port.
holddown <0-360>	Configures the RIP holddown timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.
in-policy WORD<0-64>	Configures the policy name for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when RIP adds it to the routing table.
listen enable	Specifies that the routing switch learns RIP routes through this interface. If enabled, the switch listens for a default route without listening for all routes. The default is enable.
out-policy WORD<0-64>	Configures the policy name for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement. <i>WORD<0-64></i> is a string of length 0–64 characters.
poison enable	Enables Poison Reverse. If you disable Poison Reverse (<i>no poison enable</i>), Split Horizon is enabled. By default, Split Horizon is enabled. If you enable Split Horizon, the interface does not advertise IP routes learned from an immediate neighbor back to the neighbor. If you enable Poison Reverse, the RIP updates sent to a neighbor from which a route is learned are poisoned with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network. These mechanisms prevent routing loops.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Table continues...

Variable	Value
receive version <rip1 rip2 rip1orrip2>	Indicates which RIP update version to accept on this interface. The default is rip1orrip2.
send version <notsend rip1 rip1comp rip2>	Indicates which RIP update version the router sends from this interface. ripVersion1 implies sending RIP updates that comply with RFC1058. rip1comp implies broadcasting RIP2 updates using RFC1058 route subassumption rules. The default is rip1Compatible.
supply enable	Specifies that the switch advertises RIP routes through the port. The default is enable.
timeout <15-259200>	Configures the RIP timeout interval in seconds. The default is 180.
triggered enable	Enables automatic triggered updates for RIP.

Configuring route redistribution to RIP

Configure a redistribute entry to announce certain routes into the RIP domain, including static routes, direct routes, RIP, Open Shortest Path First (OSPF), IS-IS, or Border Gateway Protocol (BGP). Optionally, use a route policy to control the redistribution of routes.

Before you begin

- Enable RIP globally.
- Configure a route policy.

Procedure

1. Enter RIP Router Configuration mode:

```
enable
configure terminal
router rip
```

2. Create the redistribution instance:

```
redistribute <bgp|direct|isis|ospf|rip|static> [vrf-src WORD<0-16>]
```

3. Apply a route policy, if required:

```
redistribute <bgp|direct|isis|ospf|rip|static> route-map WORD<0-64>
[vrf-src WORD<0-16>]
```

4. Configure other parameters.

5. Enable the redistribution:

```
redistribute <bgp|direct|isis|ospf|rip|static> enable [vrf-src
WORD<0-16>]
```

6. Ensure that the configuration is correct:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution:

```
ip rip apply redistribute <bgp|direct|isis|ospf|rip|static> [vrf
WORD<1-16>] [vrf-src WORD<0-16>]
```

Example

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router rip
Switch:1(config-rip)#redistribute rip
Switch:1(config-rip)#redistribute rip route-map test1
Switch:1(config-rip)#redistribute rip enable
Switch:1(config-rip)#exit
Switch:1(config)#ip rip apply redistribute rip
```

Variable definitions

Use the data in the following table to help you use the **redistribute** command.

Variable	Value
metric <0-65535>	Configures the metric to apply to redistributed routes.
route-map WORD<0-64>	Configures the route policy to apply to redistributed routes.
[vrf-src WORD<0-16>]	Specifies the optional source VRF instance. You can use this variable with the other command variables.
WORD<0-32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Use the data in the following table to use the **show ip rip redistribute** command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance.
vrfids WORD<1-512>	Specifies a range of VRF IDs.

Use the data in the following table to use the **ip rip apply redistribute** command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance.
vrf-src WORD<0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<bgp direct isis ospf rip static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Configuring interVRF route redistribution for RIP

Configure a redistribute entry to announce certain routes into the RIP domain, including static routes, direct routes, RIP, OSPF, IS-IS, or BGP. Use a route policy to control the redistribution of routes.

Before you begin

- Enable RIP globally.
- Configure a route policy.
- Configure the VRFs.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create the redistribution instance:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static>
```

3. Apply a route policy, if required:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static> route-map
WORD<0-64> [vrf-src WORD<0-16>]
```

4. Configure other parameters.

5. Enable the redistribution:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static> enable [vrf-
src WORD<0-16>]
```

6. Ensure that the configuration is correct:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<1-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution:

```
ip rip apply redistribute <bgp|direct|isis|ospf|rip|static> [vrf
WORD<1-16>] [vrf-src WORD<0-16>]
```

Example

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
```

```
Switch:1#router vrf red
Switch:1(router-vrf)#ip rip redistribute ospf
Switch:1(router-vrf)#ip rip redistribute ospf route-map test1
Switch:1(router-vrf)#ip rip redistribute ospf enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip rip apply redistribute ospf
```

Variable definitions

Use the data in the following table to use the **ip rip redistribute <bgp|isis|ospf|static|direct|rip>** command.

Variable	Value
<bgp direct isis ospf rip static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.
vrf-src <i>WORD</i> <0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
metric <0-65535>	Configures the metric to apply to redistributed routes.
route-map <i>WORD</i> <0-64>	Configures the route policy to apply to redistributed routes.

Use the data in the following table to use the **show ip rip redistribute** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF instance.
vrfids <i>WORD</i> <1-512>	Specifies a range of VRF IDs.

Use the data in the following table to use the **ip rip apply redistribute** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF instance.
vrf-src <i>WORD</i> <0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<bgp direct isis ospf rip static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Forcing a RIP Update for a Port or VLAN

Force RIP to update the routing table so that the port or VLAN uses the latest routing information.

About this task

If you perform this procedure, you also update the tables for all VRFs associated with the port or VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-
port]][,...]} or interface vlan <1-4059>
```

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable the triggered-update flag:

```
ip rip triggered enable
```

*** Note:**

You can enable this flag in either the GigabitEthernet or VLAN Interface Configuration mode. However, you can update the RIP routes in the GigabitEthernet Interface Configuration mode only.

3. Update the routing table:

```
action triggerRipUpdate
```

Example

Update the routing table:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#ip rip triggered enable
```

Viewing the RIP redistribution configuration information

Displays the RIP redistribution configuration information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the RIP redistribution configuration information:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the RIP redistribution configuration information:

```
Switch(config-ospf)#show ip rip redistribute
```

```
=====
                        RIP Redistribute List - GlobalRouter
=====
SRC-VRF          SRC  MET  ENABLE  RPOLICY
-----
```

```
GlobalRouter      ISIS 0      FALSE
```

Variable definitions

Use the data in the following table to use the `show ip rip redistribute` command or the `show ipv6 rip redistribute` command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

RIP configuration using EDM

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use Enterprise Device Manager (EDM) to configure and manage RIP.

Configuring RIP globally

Configure RIP global parameters on the switch so you can control RIP behavior on the system.

Before you begin

- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About this task

All router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

You can configure RIP on interfaces while RIP is globally disabled. This way, you can configure all interfaces before you enable RIP for the switch.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Globals** tab.
4. Select the **enable** option button.
5. Configure other global RIP parameters as required.
6. Click **Apply**.

Globals field descriptions

Use the data in the following table to use the Globals tab.

Name	Description
Operation	Enables or disables RIP on all interfaces. The default is disabled.
UpdateTime	Specifies the time interval between RIP updates for all interfaces. The default is 30 seconds, and the range is 0–360.
RouteChanges	Specifies the number of route changes RIP made to the IP route database. RouteChanges does not include the refresh of a route age.
Queries	Specifies the number of responses sent to RIP queries received from other systems.
HoldDownTime	Configures the length of time that RIP continues to advertise a network after the network is unreachable. The range is 0–360 seconds. The default is 120 seconds.
TimeOutInterval	Configures the RIP timeout interval. The range is 15–259200 seconds. The default is 180 seconds.
DeflImportMetric	Configures the default import metric used to import a route into a RIP domain. To announce OSPF internal routes into a RIP domain, if the policy does not specify a metric, you must use the default import metric. OSPF external routes use the external cost. The range is 0–15 and the default is 8.

Viewing RIP status

View RIP status.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Status** tab.

Status field description

Use the following table to use the RIP Status tab.

Name	Description
Address	Specifies the IP address of the router interface.
RcvBadPackets	Specifies the number of RIP response packets received by the RIP process which were subsequently discarded; for example, version 0 packet, or an unknown command type.

Table continues...

Name	Description
RcvBadRoutes	Specifies the number of routes, in valid RIP packets, that are ignored; for example, unknown address family, or invalid metric.
SentUpdates	Specifies the number of triggered RIP updates sent on this interface, that do not include full updates sent containing new information.

Configuring RIP interface compatibility

Configure RIP parameters on an interface so you can control RIP behavior on the interface. You can specify the RIP version to use on interfaces that you configure to send (supply) or receive (listen to) RIP updates.

Before you begin

- Configure a routing interface (either a router port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIP globally.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About this task

On an interface, RIP does not operate until you enable it globally and on the interface.

Although visible, the switch does not support the AuthType and AuthKey parameters.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Interface** tab.
4. Double-click the **Send** value to edit it, and then select the RIP version datagrams the router sends.
5. Double-click the **Receive** value to edit it, and then select the RIP version datagrams for which the router listens.
6. Click **Apply**.

Interface Field Descriptions

Use the data in the following table to use the Interface tab.

Name	Description
Address	Specifies the IP address of the router interface.

Table continues...

Name	Description
Domain	Specifies the value inserted into the routing domain parameter of all RIP packets sent on this interface.
AuthType	Specifies the type of authentication to use on this interface.
AuthKey	Specifies the authentication key whenever AuthType is not noAuthentication.
Send	Specifies the update version the router sends on this interface: • DoNotSend—no RIP updates sent on this interface • ripVersion1—RIP updates compliant with RFC1058 • rip1Compatible—broadcast RIPv2 updates using RFC1058 route subassumption rules • ripVersion2—multicast RIPv2 updates The default is rip1compatible.
Receive	Indicates which versions of RIP updates to accept: • rip1 • rip2 • rip1OrRip2 The default is rip1OrRip2. Rip2 and rip1OrRip2 imply receipt of multicast packets.

Job Aid

Choose one of three options for receiving RIP updates:

- rip1OrRip2—accepts RIPv1 or RIPv2 updates
- rip1—accepts RIPv1 updates only
- rip2—accepts RIPv2 updates only

The following table describes the four RIP send modes that the switch supports. You can configure RIP send modes on all router interfaces.

Table 7: RIP send modes

Send mode	Description	Result
rip1Compatible	Broadcasts RIPv2 updates using RFC1058 route consumption rules. This is the default mode.	• Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.0.2.255) • RIP update is formed as a RIP-2 update, including network mask • RIP version = 2

Table continues...

Send mode	Description	Result
ripVersion1	Broadcasts RIP updates that are compliant with RFC1058	• Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 198.0.2.255) • RIP update is formed as a RIP-1 update, no network mask included • RIP version = 1
ripVersion2	Broadcasts multicast RIPv2 updates	• Destination MAC is a multicast, 01-00-5e-00-00-09 • Destination IP is the RIP-2 multicast address, 224.0.0.9 • RIP update is formed as a RIP-2 update including network mask • RIP version = 2
doNotSend	Does not send RIP updates on the interface	None

Configuring RIP on an interface

Configure RIP parameters to control and optimize RIP routing for the interface.

Before you begin

- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Interface Advance** tab.
4. Double-click a RIP parameter to edit it, as required.
5. Click **Apply**.

Interface Advance field descriptions

Use the data in the following table to use the RIP Interface Advance tab.

Name	Description
Address	Shows the address of the entry in the IP RIP interface table.
Interface	Indicates the index of the RIP interface.
Enable	Shows if the RIP interface is enabled or disabled.

Table continues...

Name	Description
Supply	Enables (true) or disables (false) the ability to send RIP updates on this interface.
Listen	Configures whether the switch learns routes on this interface.
Poison	Configures whether to advertise RIP routes learned from a neighbor back to the neighbor. If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, RIP poisons the RIP updates, sent to the neighbor from which a route is learned, with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network.
DefaultSupply	Enables (true) or disables (false) an advertisement of a default route on this interface. This command takes effect only if a default route exists in the routing table.
DefaultListen	Enables (true) or disables (false) the switch to accept the default route learned through RIP on this interface. The default is disabled. Enable DefaultListen to add a default route to the route table if another route advertises it.
TriggeredUpdate	Enables (true) or disables (false) the switch to send RIP updates from this interface.
AutoAggregate	Enables (true) or disables (false) automatic route aggregation on this interface. If enabled, the switch automatically aggregates routes to their natural mask when an interface advertises them. The default is disabled.
InPolicy	Determines if RIP can learn routes on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if RIP advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The range is 1–15. The default is 1.

Job aid

The following table indicates the relationship between switch action and the RIP supply and listen settings.

Table 8: RIP supply and listen settings and switch action

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
Disabled (false)	Disabled (false)			Sends no RIP updates.

Table continues...

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
Enabled (true)	Disabled (false)			Sends RIP updates except the default.
Disabled (false)	Disabled (false)			Sends only the default (default route must exist in routing table).
Enabled (true)	Enabled (true)			Sends RIP updates including the default route (if it exists).
		Disabled (false)	Disabled (false)	Does not listen to RIP updates.
		Enabled (true)	Disabled (false)	Listens to all RIP updates except the default.
		Disabled (false)	Enabled (true)	Listens only to the default.
		Enabled (true)	Enabled (true)	Listens to RIP updates including the default route (if it exists).

Configuring RIP on a port

Configure RIP on a port so that the port can participate in RIP routing.

Before you begin

- Assign an IP address to the port.
- Configure RIP and enable it globally.

Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

- Enable RIP on the interface.

About this task

On an interface, RIP does not operate until you enable it globally and on the interface.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
3. Click **IP**.
4. Click the **RIP** tab.
5. Configure the RIP parameters as required.
6. Click **Apply**.

RIP field descriptions

Use the data in the following table to use the RIP tab.

Name	Description
Enable	Enables or disables RIP on the port.
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the port is down. The default is false. If you configure a port with no link and enable AdvertiseWhenDown, the port does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if this interface advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.

Table continues...

Name	Description
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

Configuring RIP on a VLAN

Configure RIP on a VLAN so that the VLAN acts as a routed VLAN (a virtual router).

Before you begin

- Configure the VLAN.
- Assign an IP address to the VLAN.
- Enable RIP globally.
- Enable RIP on the interface.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **RIP** tab.
7. Configure the VLAN RIP parameters as required.
8. Click **Apply**.

RIP field descriptions

Use the data in the following table to use the RIP tab.

Name	Description
Enable	Enables or disables RIP on the VLAN.
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.

Table continues...

Name	Description
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the interface is down. The default is false. If you configure a VLAN with no link and enable AdvertiseWhenDown, the VLAN does not advertise the route until the VLAN is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if this interface advertises a route from the routing table. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

Configuring interVRF route redistribution policies

Configure interVRF route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF, RIP, or BGP. Use a route policy to control the redistribution of routes.

Before you begin

- VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **Policy**.
3. Click the **Route Redistribution** tab.
4. Click **Insert**.
5. Click the ellipsis (...) button near the **DstVrflid** box to select the source and destination VRF IDs.
6. Click the ellipsis (...) button near the **SrcVrflid** box to select the source and destination VRF IDs.
7. In the **Protocol** option box, select the protocol.
8. In the **RouteSource** option box, select the route source.
9. Select **enable**.
10. Click the ellipsis (...) button near the **RoutePolicy** box to choose the route policy to apply to the redistributed routes.
11. Configure other parameters as required.
12. Click **Insert**.
13. Click the **Applying Policy** tab.
14. Select **RedistributeApply**.
15. Click **Apply**.

Route Redistribution field descriptions

Use the data in the following table to use the Route Redistribution tab.

Name	Description
DstVrflid	Specifies the destination VRF ID to use in the redistribution.

Table continues...

Name	Description
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrfId	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
Subnets	Indicates that subnets must be advertised individually (applies to OSPF only).

Configuring route redistribution to RIP

Configure a redistribute entry to announce routes of a certain source protocol type into the RIP domain, for example, static, RIP, or direct. Use a route policy to control the redistribution of routes.

Before you begin

- Enable RIP globally.
- Configure a route policy.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About this task

Important:

Changing the RIP redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. It is recommended that if you want to change default preferences for a RIP redistribute context, you must do so before you enable the protocols.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Redistribute** tab.
4. Click **Insert**.

5. Configure the source of the routes to redistribute.
6. Select **enable**.
7. Select the route policy to apply to redistributed routes.
8. Configure a metric value.
9. Click **Insert**.

Redistribute field descriptions

Use the data in the following table to use the Redistribute tab.

Name	Description
DstVrflid	Specifies the destination VRF instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrflid	Specifies the source VRF instance. You cannot configure this variable.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) a RIP redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) that redistributes external routes from a specified source into an RIP domain. Click the ellipsis (...) button and choose from the list in the Route Policy dialog box.
Metric	Configures the RIP route redistribution metric for basic redistribution. The value can be in the range 0–65535. A value of 0 indicates to use the original cost of the route.

Glossary

area border router (ABR)	A router attached to two or more areas inside an Open Shortest Path First (OSPF) network. Area border routers play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the Autonomous System, and using an EGP to route packets to other Autonomous Systems.
autonomous system border router (ASBR)	A router attached at the edge of an OSPF network. An ASBR uses one or more interfaces that run an interdomain routing protocol such as BGP. In addition, a router distributing static routes or Routing Information Protocol (RIP) routes into OSPF is considered an ASBR.
backup designated router (BDR)	A router that assumes the designated router (DR) role for the Open Shortest Path First (OSPF) protocol if the DR fails.
Circuitless IP (CLIP)	A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface.
classless interdomain routing (CIDR)	The protocol defined in RFCs 1517 and 1518 for using subnetwork masks, other than the defaults for IP address classes.
database description (DD) packets	Exchanged when a link is initially established between neighboring routers that synchronizes their link state databases. The Open Shortest Path First (OSPF) protocol uses DD packets.
designated router (DR)	A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.

equal cost multipath (ECMP)	Distributes routing traffic among multiple equal-cost routes.
Interior Gateway Protocol (IGP)	Distributes routing information between routers that belong to a single Autonomous System (AS).
Internal Router (IR)	A router with interfaces only within a single area inside an Open Shortest Path First (OSPF) network.
Internet Protocol Control Packet (IPCP)	Establishes and configures Internet Protocol data transmission over a Point-to-Point Protocol link.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
link-state advertisement (LSA)	Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.
link-state database (LSDB)	A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the autonomous system (AS), with itself at the root of each path.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
Message Digest 5 (MD5)	A one-way hash function that creates a message digest for digital signatures.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
nonbroadcast multiaccess (NBMA)	Interconnects multiple devices over a broadcast network through point-to-point links. NBMA reduces the number of IP addresses required for point-to-point connections.
not so stubby area (NSSA)	Prevents the flooding of external link-state advertisements (LSA) into the area by providing them with a default route. An NSSA is a configuration of the Open Shortest Path First (OSPF) protocol.

Open Shortest Path First (OSPF)	A link-state routing protocol used as an Interior Gateway Protocol (IGP).
prefix	A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.
remote monitoring (RMON)	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
shortest path first (SPF)	A class of routing protocols that use Dijkstra's algorithm to compute the shortest path through a network, according to specified metrics, for efficient transmission of packet data.
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
type of service (TOS)	A field in the IPv4 header that determines the Class of Service prior to the standardization of Differentiated Services.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
variable-length subnet masking (VLSM)	Allocating IP addressing resources to subnets according to their individual need rather than some general network-wide rule.