

Configuring IPv6 Routing for VOSS

© 2017-2020, 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	
Purpose	
Conventions	10
Text Conventions	10
Documentation and Training	12
Getting Help	12
Providing Feedback	13
Chapter 2: New in this Document	14
Notice about Feature Support	14
Chapter 3: IPv6 basics	15
Origins of IPv6	15
Advantages of IPv6	15
Comparison of IPv4 and IPv6	16
IPv6 packet	17
IPv6 header	17
IPv6 extension headers	18
IPv6 address component summary	20
IPv6 address formats	21
Address types	21
Unicast addresses	21
Multicast addresses	23
Anycast	25
IP address prefix	26
Interface ID	26
How to write an IPv6 address	27
ICMPv6	27
Path MTU discovery	28
Routing	29
Route scaling	33
IPv6 Circuitless IP	34
Equal Cost Multipath	
ECMP with static routes	35
Disable IPv6 ICMP multicast	36
Route Policy Definition	36
IPv6 Basic Configuration using CLI	41
Enabling the IPv6–mode boot config flag	
Configuring an IPv6 static neighbor address	42
Configuring an IPv6 interface	
Assigning IPv6 addresses to a brouter port or VLAN	49

Configuring IPv6 route preferences	50
Viewing global IPv6 information	52
Creating IPv6 static routes	56
View Routes Information	58
Creating an IPv6 CLIP interface	60
Enabling IPv6 ECMP	61
Configuring maximum number of ECMP paths	62
Enabling or disabling IPv6 ICMP multicast	63
Enabling Stateless Address Autoconfiguration	64
Configuring route advertisement on the management port	64
Configuring process-redirect for management port	66
Viewing IPv6 default routers	67
Configuring an IPv6 prefix list	68
Configuring IP Route Policies	69
IPv6 Basic Configuration using EDM	76
Enabling the IPv6–mode boot config flag	76
Configuring IPv6 globally	77
Configuring an IPv6 interface	
Configuring an IPv6 brouter port interface	81
Configuring an IPv6 VLAN interface	
Assigning IPv6 addresses to interfaces	
Assigning IPv6 addresses to a brouter port interface	87
Assigning an IPv6 address to a VLAN	89
Creating IPv6 static routes	
Configuring IPv6 route preferences	
View Route Information	92
Viewing IPv6 Default Routers	
Configure a Circuitless IPv6 Interface	
Configuring IPv6 Prefix List	
Configuring an IPv6 Route Policy	
Configuring IPv6 Route Redistribution Policies	100
Chapter 4: Neighbor discovery	. 102
Neighbor discovery	. 102
ND messages	104
Neighbor discovery cache	
Router discovery	106
Host autoconfiguration	
Neighbor Discovery Configuration using CLI	
Configuring an IPv6 discovery prefix	
Configuring route advertisement	. 110
Configuring the neighbor cache	116
Viewing cached destination information	
Neighbor Discovery Configuration using EDM	. 119

	Configuring an IPv6 discovery prefix	119
	Configuring an IPv6 discovery prefix port	121
	Configuring an IPv6 discovery prefix on a VLAN	123
	Configuring route advertisement	
	Configuring route advertisement on an IPv6 interface for a brouter port	128
	Configuring route advertisement on an IPv6 interface for a VLAN	130
	Configuring the neighbor cache	
	Viewing cached destination information	133
Ch	apter 5: DHCP Relay	135
	DHCP Relay	
	DHCP Relay Network Topology and Workflow	
	DHCP Relay Configuration using CLI.	
	Configuring a DHCP Relay forwarding path	
	Configuring DHCP Relay for an interface	
	Viewing DHCP Relay information	
	DHCP Relay Configuration using EDM	
	Configuring a DHCP Relay forwarding path	
	Configuring DHCP Relay for an interface	
	Modifying DHCP Relay for a VLAN	
	Modifying DHCP Relay for a port	
Ch	apter 6: Tunneling	
	Tunneling	
	Manually configured tunnels	
	Limitations	
	Tunneling Configuration using CLI	
	Configuring a tunnel	
	Viewing tunnel interfaces	
	Modifying tunnel hop limits	
	Tunneling Configuration using EDM	
	Configuring a tunnel	
	Modifying tunnel hop limits	153
	Modifying tunnel hop limits for a specific tunnel	
	Viewing IPv6 addresses on a tunnel	
Ch	apter 7: OSPFv3	
	OSPFv3	
	IPsec support with OSPFv3	
	OSPF Graceful Restart	
	Helper Mode	
	ECMP with OSPFv3	
	OSPFv3 and Route Redistribution	
	OSPFv3 Configuration using CLI	
	Configuring OSPF globally	
	Creating an OSPF area	166

Creating OSPF area ranges	168
Creating an OSPF virtual link	169
Configuring IPsec for the OSPF virtual link	171
Configuring OSPF default metrics	173
Configuring OSPF on a port or VLAN	175
Configuring OSPF on a tunnel	178
Viewing OSPFv3 Information	181
Viewing OSPFv3 Default Cost Information	184
Adding an NBMA neighbor	185
Configuring link LSA suppression	186
Configuring Route Redistribution to OSPFv3 in GRT mode	187
Viewing the Status of OSPFv3 Redistribution	
Disabling Helper mode for OSPFv3	190
OSPFv3 Configuration using EDM	191
Configuring OSPFv3 globally	191
Creating an OSPFv3 Area	193
Creating OSPFv3 Area Ranges	195
Creating an OSPFv3 Virtual Link	196
Configuring IPsec for the OSPF virtual link	198
Creating an OSPF interface on a brouter port	200
Create an OSPF VLAN Interface	202
Creating an OSPF interface on a tunnel	205
Viewing the AS-scope link-state database	207
Viewing the area-scope LSDB	208
Viewing the link-scope LSDB	209
Adding an NBMA neighbor	210
Configuring Route Redistribution to OSPFv3	212
Modifying an OSPFv3 interface	213
Viewing OSPFv3 neighbors	215
Viewing virtual neighbors	217
Chapter 8: RIPng	219
RIPng fundamentals	219
RIPng Configuration using CLI	221
Configuring RIPng globally	221
Configuring RIPng on an interface	222
Configuring RIPng custom values	223
Configuring RIPng route distribution	225
RIPng Configuration using EDM	226
Configuring RIPng globally	226
Configuring an IPv6 RIPng interface	227
Configuring an IPv6 RIPng VLAN interface	228
Configuring an IPv6 RIPng brouter port interface	230
Graphing IPv6 RIPng statistics	231

Configuring route redistribution to RIPng	232
Viewing stats for RIPng interfaces	
Chapter 9: VRRP	234
· VRRP	
VRRPv3	
VRRPv3 guidelines	239
VRRP Configuration using CLI	
Configuring the VRRP interface	
Viewing VRRP information	
Configuring VRRP notification control	
Configuring additional VRRP parameters for an interface	
Enabling IPv6 VRRP preempt-mode	
VRRP Configuration using EDM	250
Configure VRRP for an Interface	
Configuring VRRP notification control	252
Configuring additional addresses on the VRRP brouter port	253
Configuring additional addresses on the VRRP interface	254
Chapter 10: RSMLT	255
RSMLT	
RSMLT Configuration using CLI	
Configuring RSMLT on a VLAN	
Enabling RSMLT Edge support	258
Viewing RSMLT information	259
RSMLT Configuration using EDM	260
Configuring RSMLT on a VLAN	261
Enabling RSMLT Edge support	262
Modifying the RSMLT local information	
Modifying RSMLT peer information	263
Viewing RSMLT Edge peers	264
Chapter 11: Viewing IPv6 connections	266
Viewing IPv6 Connections using CLI	
Viewing TCP and UDP information	
Viewing IPv6 Connections using EDM	
Viewing TCP global information	268
Viewing TCP connections information	269
Viewing TCP listeners information	270
Viewing UDP endpoint information	271
Chapter 12: IPv6 Alternative Routes	273
Alternative routes	
Enable IPv6 Alternative Routes	
Chapter 13: IPv6 configuration examples	
IPv6 tunnels	
	280

Contents

282
294
295
297
299
300
301
301
304
305
307
307
308
308
309
310
311
312
313
315

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 (includes VSP 4450 Series)
- Extreme Networks VSP 4900 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

Note:

VOSS is licensed on the XA1400 Series as a Fabric Connect VPN (FCVPN) application, which includes a subset of VOSS features. FCVPN transparently extends Fabric Connect services over third-party provider networks.

This document provides conceptual and procedural information to configure IPv6 routing operations. Included in this document are Operations, Administration, and Management (OA&M), DHCP Relay, Virtual Router Redundancy Protocol (VRRP), static routes, Open Shortest Path First version 3 (OSPFv3), IPv6 tunnels, and Routed Split MultiLink Trunking (RSMLT).

Note:

The software does not support IPv4-mapped IPv6 addresses, for example, 0::FFFF:a.b.c.d, or IPv4-compatible IPv6 addresses, for example, 0::a.b.c.d.

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.

IPv6 uses the following key security features: SNMP version 3 (SNMPv3) and Secure Shell (SSH).

For detailed information about SNMPv3, see Configuring Security for VOSS.

For detailed information about SSH, see Administering VOSS.

For information about IPv6 shortcuts and IPv6 Inter-VSN routing, see <u>Configuring Fabric Basics and Layer 2 Services for VOSS</u>.

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 cfm maintenance-domain maintenance-level <0-7>, you can enter cfm maintenance-domain maintenance-level 4.	
Bold text	Bold text indicates the GUI object name you must act upon.	

Table continues...

Convention	Description
	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 ip address {A.B.C.D}, 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 show clock [detail], you can enter either show clock or show clock detail.
Ellipses ()	An ellipsis () indicates that you repeat the last element of the command as needed.
	For example, if the command syntax is ethernet/2/1 [<parameter> <value>], you enter ethernet/2/1 and as many parameter-value pairs as you need.</value></parameter>
Italic Text	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:
	• show ip route
	• Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]
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

Table continues...

Convention	Description	
	not type the vertical line when you enter the command.	
	For example, if the command syntax is access- policy by-mac action { allow deny }, you enter either access-policy by-mac action allow Or access-policy by-mac action deny, but not both.	

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation

Release Notes

Hardware/software compatibility matrices for Campus and Edge products

Supported transceivers and cables for Data Center products

Other resources, like white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, 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: (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, or serial numbers for all involved Extreme Networks products

- · A description of the failure
- A description of any actions 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

Subscribe 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 (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. Select Submit.

Providing Feedback

The Information Development team at Extreme Networks has 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 we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

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 physical hardware restrictions, see your hardware documentation.

Chapter 3: IPv6 basics

This chapter provides concepts and procedures to complete basic IPv6 configuration, for example, IPv6 forwarding and static routes.

Origins of IPv6

The growth of IP address use is exponential.

Predictions indicated that the IPv4 address pool could be exhausted as early as 1994.

So, in July 1991, the Internet Engineering Task Force (IETF) began researching a replacement for IPv4.

That replacement is IPv6.

The Internet Assigned Numbers Authority (IANA) free pool of IPv4 addresses reached 0% in February 2011, according to the American Registry for Internet Numbers (ARIN).

While IPv4 addresses may remain available for some time within reserved pools, no further IPv4 addresses are available for reservation.

Although IPv6 is designed to replace IPv4, IPv6 is not backward-compatible and IPv4 and IPv6 need to coexist within your network during and after the transition to IPv6.

Advantages of IPv6

IPv6 can provide more addresses and support more networks than IPv4. For example, IPv6 offers enough addresses for every person on Earth to have 1 million addresses.

Because IPv6 offers a larger address space it offers improved scalability.

Following are additional advantages of IPv6 over IPv4:

- With 128 bit addresses, the larger IPv6 address space offers global access and scalability and solves the pending exhaustion of IP addresses.
- Network Address Translation (NAT) is no longer required.

Flat address space and transparency are restored by IPv6 because NAT is eliminated.

- Routing efficiency is improved due to the hierarchical network architecture.
 - IPv6 allows for hierarchical routing and effective route summarization.
- IPv6 supports Auto-configuration.
- IPv6 supports plug-and-play.
- Enhanced support is included for mobile IP and mobile computing devices.
 - Addresses can be permanently assigned to end devices such as DSL, PDAs, mobile terminals and PCs.
- Neighbor discovery (ND) replaces ARP in IPv6.
 - ND combines the IPv4 services for IPv4 Address Resolution Protocol (ARP) and router discovery.

Comparison of IPv4 and IPv6

The following table compares the key differences between IPv4 and IPv6.



This information may not reflect IPv6 support in the current release.

Table 3: IPv4 and IPv6 key differences compared

Feature	IPv4	IPv6
Address length	32 bits	128 bits
IPsec support	Optional	Required
QoS support	Limited	Improved
Fragmentation	Hosts and routers	Hosts only
Minimum MTU packet size	576 bytes	1280 bytes
Checksum in header	Yes	No
Options in header	Yes	Yes
Link-layer address resolution	ARP (broadcast)	Multicast neighbor discovery messages
Multicast membership	IGMP	Multicast Listener Discovery
Router discovery	Optional	Optional
Uses broadcasts	Yes	No
Address configuration	Manual, DHCP	Automatic, DHCP

IPv6 packet

Each IPv6 packet can include mandatory and non-mandatory components.

An IPv6 packet includes:

- The basic header, which has a fixed length and is mandatory
- Extension header(s), which has a variable length and is not mandatory
- Payload, which has a variable length and is not mandatory

The following figure illustrates the components of an IPv6 packet.

Basic header	Extension header(s)	Payload
--------------	---------------------	---------

Figure 1: IPv6 packet components



Nodes must be able to handle packets up to 1,280 octets in length.

IPv6 header

The IPv6 header basic length is fixed at 40 octets (bytes) and it contains the following fields:

Table 4: Fields in the IPv6 header

Field	Size in bits
Ver—Internet Protocol version number, with a value of 6	4
DS byte—Traffic class field, similar to Type of Service in IPv4	8
Flow label—identifies traffic flow for additional Quality of Service (QoS)	20
Payload Length—Unsigned integer, the length of the IPv6 payload	16
Next header selector—identifies the next header	8
Hop limit unsigned integer—decrements by 1 each time a node forwards the packet (nodes discard packets with hop limit values of 0)	
Source address	128
Destination address	128

The following figure illustrates the basic IPv6 header, without extension headers.

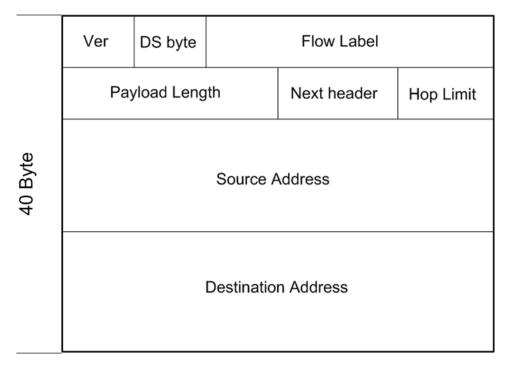


Figure 2: IPv6 header

IPv6 extension headers

IPv6 extension headers describe processing options.

Each extension header contains a separate category of options and is identified by a number, similar to protocol identification numbers.

An IPv6 packet can include extension headers, but they are not mandatory.

The following figure illustrates the IPv6 header with extension headers.

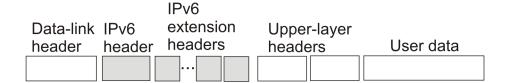


Figure 3: IPv6 header with extension headers

IPv6 examines the destination address in the main header of each packet it receives.

This examination determines whether the router is

• the packet destination - if the router is the packet destination, IPv6 examines the header extensions that contain options for destination processing.

• an intermediate node in the packet data path - if the router is an intermediate node, IPv6 examines the header extensions that contain forwarding options.

By examining only the extension headers that apply to the operations it performs, IPv6 reduces the amount of time and resources required to process a packet.

IPv6 defines the following extension headers as described in the following table:

Table 5: IPv6 extension headers

Extension header name	Description	
hop-by-hop	Contains optional information, and sub-options for Router Alert and Jumbo Payload, that all intermediate IPv6 routers examine between the source and the destination.	
destinations-options	Contains optional information for the destination node.	
	This option can appear twice, once for way points and once for final destination.	
source-routing	Contains a list of one or more intermediate nodes that define a path for the packet to follow through the network to the destination.	
	The packet source creates this list.	
	The source-routing function is similar to the IPv6 source routing options.	
fragmentation	Uses an IPv6 source to send packets larger than the size specified for the path maximum transmission unit (MTU).	
authentication	provides security for IPv6 datagrams	
encapsulated security payload (ESP)	provides security for IPv6 datagrams	
The authentication extension header and the encapsulated security payload extension header can be used		

together to provide security services for IPv6 datagrams.

The recommended extension header order is:

- Hop-by-hop
- Destination option 1
- Routing
- Fragmentation
- Authentication/ESP
- Destination Option 2

The presence of particular extension headers within a packet can cause slower packet processing if the IPv6 implementation handles only certain headers and diverts others to a slow path. For example, many IPv6 implementations usually process Hop-by-Hop extension headers on the control plane.

IPv6 address component summary

The IPv6 Internet is divided into addressing zones and IPv6 addresses can be categorized by type and scope.

IPv6 addressing is represented in RFC 4291.

Address types

IPv6 addresses are divided into the following types:

Unicast

Unicast addresses provide one-to-one communication.

Multicast

Multicast addresses are similar in operation to IPv4 and provide one-to-many communication.

Anycast

An Anycast address is a Unicast address used for several devices to allow them to communicate with the device closest to the source; one-to-nearest communication.

Broadcast

In IPv6, broadcast addresses have been superseded by multicast addresses per RFC 4291.

For more information about address types and scopes, see IPv6 Address Types on page 21.

Address scopes

Following are IPv6 address scopes:

- node-local
- link-local
- global

The switch does not support site-local addresses and, according to RFC 4193, site-local addresses will be replaced by unique-local addresses.

For more information about address types and scopes, see IPv6 address formats on page 21.

Address zones

The IPv6 Internet is divided into zones.

For example:

- Each node is a separate zone of the node-local scope.
- Each link is a separate zone of the link-local scope.
- The entire Internet is a single zone of global scope.

Zones of the same scope do not overlap.

IPv6 address formats

IPv6 addresses are 128 bits long. In comparison, IPv4 addresses are 32 bits in length.

The IPv6 address contains an

- · address type
- · address prefix
- · interface ID

The following figure illustrates the IPv6 address format.

Type Address prefix	Interface ID
---------------------	--------------

Figure 4: IPv6 address format

Address types

IPv6 uses three main address types to help route packets.

Address types are:

- Unicast: global, link—local, special unspecified, special loopback
- Multicast
- Anycast

Difference between multicast and anycast

Anycast address delivery is from one to one-of-many, whereas multicast address delivery is from one to many.

Unicast addresses

Unicast addresses provide one-to-one communication

Unicast addresses provide one-to-one communication.

Global

A Unicast global address identifies a single interface and is similar to an IPv4 public address.

Unicast global addresses are globally routable in the same manner as IPv4 addresses.

The following figure illustrates the Unicast global address parts.

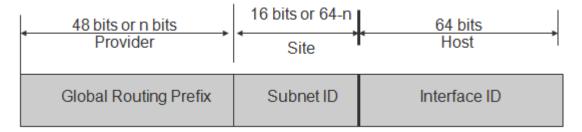


Figure 5: Unicast global address parts

An IPv6 Unicast global address is composed of the following 3 levels:

- public topology (48 bit Global Routing Prefix)
 - 001, specifies an IPv6 Unicast global address
 - Top Level Aggregation Identifier (TLA ID), the highest level in routing hierarchy
 - Res, reserved for future use
 - Next Level Aggregation Identifier (NLA ID), specifies a customer site
- site topology (16 bit Subnet ID)
 - Site Level Aggregation Identifier (SLA ID); assigned within the site, an ISP cannot affect the SLA ID, enables up to 65,536 subnets within a site
- interface ID (64 bits)
 - specifies the interface for a node on a subnet

The system uses the 48 bit global routing prefix for the route prefix exchange.

The IPv6 Prefix for Unicast global is 2000::/3 (RFC3513).

Link-local

Hosts on the same link/subnet use automatically configured IPv6 Unicast link-local addresses to communicate with each other.

Link-local addresses are automatically configured on all interfaces.

Routers do not forward packets containing a destination or source address with a link-local address.

IPv6 uses neighbor discovery (ND) for address resolution.

The IPv6 prefix for link-local Unicast addresses is FE80::/10 (RFC3513).

The following figure illustrates the parts of a Unicast Link-local address.

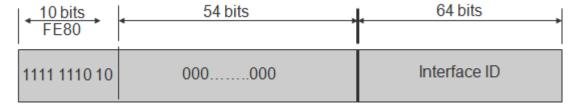


Figure 6: Unicast Link-local address

Special addresses

The Unicast/special/unspecified address indicates the absence of an address and is the only valid SRC address for IPv6 Duplicate Address Detection (DAD).

Note:

Do not assign an unspecified address, either statically or dynamically, to an interface.

The Unicast/special/loopback address is a special case Unicast address only found inside a single node.

The switch does not support the loopback address.

Equivalent to the IPv4 loopback address 127.0.0.1, represented as 0:0:0:0:0:0:0:0:0:1 or ::1; a node uses a loopback address to send a packet to itself.

The loopback address is beneficial in troubleshooting and testing the IP stack because you can use it to send a packet to the protocol stack without sending it onto the subnet.

Note:

Do not assign a loopback address, either statically or dynamically, to an interface.

Both Loopback and Unspecified addresses are not valid destination addresses.

An example of a unicast IPv6 address is 1080:0:0:0:8:8000:200C:417A

An example of a link-local Unicast IPv6 address is FE80::4445:4eff:fe54:1212

Multicast addresses

Multicast addresses provide one-to-many communication

Multicast addresses provide one-to-many communication.

An IPv6 multicast address identifies a group of nodes.

The scope is built into the multicast address structure.

The system uses a multicast address to send traffic to multiple destinations. In this situation traffic experiences less delay with a multicast address than it would with Unicast address.

The following figure shows the format of an IPv6 multicast address.

8 bits	4 bits	4 bits	112 bits
11111111	flags	scope	group ID

Figure 7: IPv6 multicast address format

A value of FF (11111111) in the 8 high-order bits of an IPv6 address indicates that the address is an IP multicast address.

The Multicast IPv6 Prefix is FF00::/8 (RFC3513).

Flags

The 4-bit flags field indicates whether the group is permanent or transient. The first 3 bits are reserved and the 4th bit represents the Transient flag. Currently only the Transient (T) flag is defined. A T flag set to 0 specifies a permanently assigned multicast address. A T flag set to 1 specifies a transient address.

Group ID

The 112 bit group ID identifies the multicast group.

An example of a multicast address is FF01:0:0:0:0:0:0:101

Scope field

The 4-bit scope field within the group ID specifies the multicast traffic scope.

Following is a list of the scope options that limit the scope of the multicast address:

- 1 node-local
- 2 link-local
- 3 subnet local
- 4 admin local
- 5 site-local not supported
- 8 organization-local
- B community-local
- E global

Examples of multicast addresses

All-nodes addresses look like this:

FF01::1 (Node Local), FF02::1 (Link Local)

All-routers addresses look like this:

FF01::2 (Node Local), FF02::2 (Link Local)

A solicited node or host address looks like this:

FF02::1:FF1E:8329.

In this case the MAC is 00-02-B3-1E-83-29 and the IPv6 address is fe80::202:B3FF:FE1E:8329.

The following table lists some well-known multicast IPv6 addresses

Table 6: Well-known multicast IPv6 addresses

Name	Address
All Nodes	FF02:0:0:0:0:0:1
All Routers	FF02:0:0:0:0:0:2

Table continues...

Name	Address
OSPFIGP	FF02:0:0:0:0:0:5
OSPFIGP Designated Routers	FF02:0:0:0:0:0:6
All PIM Routers	FF02:0:0:0:0:0:D
VRRP	FF02:0:0:0:0:0:12
All MLDv2–capable routers	FF02:0:0:0:0:0:16
All DHCP agents	FF02:0:0:0:0:0:0:2
Solicited Node address	FF02::1:FF00:0000/104

Anycast

Anycast addresses provide one-to-nearest (one to one-of-many) communication.

Anycast addresses provide one-to-nearest (one to one-of-many) communication.

An anycast address designates a set of interfaces that share an address.

A packet sent to an anycast address goes only to the nearest member of the group. Considering routing distance, the system delivers packets with anycast addresses only to the nearest member of a group of multiple interfaces.

Restrictions

An anycast address must not be:

- · used as the source address in an IPv6 packet
- assigned to an IPv6 host (you can assign an anycast address to an IPv6 router)

Anycast address scopes

Anycast addresses have the following scopes:

- · Link-local—the local link; nodes on the same subnet
- · Global—IPv6 Internet addresses

Similar to anycast IPv4 addresses, IPv6 anycast addresses are more efficient. They use the unicast address space but identify multiple interfaces.

IPv6 delivers a packet bearing an anycast address to the nearest interface identified by the address.

Currently anycast addresses are assigned to routers and are used as destination addresses. Because packets bearing anycast addresses are delivered to the closest router, you can also access the closest name server or time server with an anycast address.

Visually there is no distinction between an anycast address and a unicast address.

Note:

The switch supports only the subnet-router anycast address.

You cannot configure any specific anycast addresses beyond the automatic, generic subnet-router anycast address.

IP address prefix

Address prefixes represent one of the following:

- network identifier
- · fixed address part

Examples of IP address prefixes

2001:10F2::/48 represents a summarized route prefix

2001:10F2:0:102F::/64 represents a subnet or link prefix

FF00::/8 represents Multicast IPv6

Interface ID

Interface identifiers identify interfaces on a link.

As long as the interfaces are attached to different subnets, you can use the same identifier on more than one interface on a single node.

The IPv6 interface ID is as unique as the MAC address.

The interface ID is derived by a formula that uses the link layer 48-bit MAC address. In most cases, the interface ID is a 64-bit interface ID that contains the 48-bit MAC address. If you enter less than 64 bits, the system adds leading zeroes to extend the interface ID length to 64 bits.

You can configure the interface ID in the following ways:

- · Manual configuration
- DHCPv6 (can configure the whole address)
- Automatic derivation from EUI-64 (MAC address or other HW serial)—enables serverless or stateless auto-configuration when combined with high order part of address learned from router advertisements
- Pseudo-random generation (client privacy)—enables serverless or stateless auto-configuration when combined with high order part of address learned from router advertisements

The switch supports manual interface ID configuration or automatic derivation from EUI-64.

Note:

You must manually specify the network prefix, regardless of the interface ID formation method.

For stateless autoconfiguration, the ID is 64 bits in length.

For more information about stateless autoconfiguration, see <u>Host autoconfiguration</u> on page 106.

How to write an IPv6 address

The appearance of IPv6 addresses differs from IPv4 addresses and you express them differently.

Hexadecimal IPv6 address representations

The 128 bits in an IPv6 address are divided into 8 blocks of 16 bits each.

Following is the preferred IPv6 address format:

```
XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX
```

Each 16 bit block in an IPv6 address is converted into a 1 to 4 digit hexadecimal number separated by colons (:).

The format to represent an IPv6 address is n:n:n:n:n:n:n:n;n, where n is the hexadecimal representation of 16 bits in the address; for example, 2001:0:0:0:0:0:0:0:43.

Each nonzero field must contain at least one numeral.

Within a hexadecimal field, you do not need leading zeros.

Certain classes of IPv6 addresses commonly include multiple adjacent fields that contain hexadecimal 0.

The sample address—2001::43—includes six adjacent fields that contain zeroes represented by a double colon (::) .

You can use a double colon to compress the leading zero fields in a hexadecimal address.

A double colon can appear only once in an address.

Four more ways to write an IPv6 address

```
2001:DB8:0000:0000:25AB:0000:0000:0001
2001:DB8:0:0:25AB:0:0:1
2001:DB8:0:0:25AB::1
2001:DB8::25AB:0:0:1
```

ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) maintains and improves on features from ICMP for IPv4.

ICMPv6 reports the delivery of forwarding errors.

For example:

- Destination unreachable
- Packet too big (path MTU)
- Time exceeded (fragmentation)
- · Parameter problem

ICMPv6 also delivers information messages such as ping, otherwise known as

- Echo request
- · Echo reply

! Important:

By providing a framework for informational messages, ICMPv6 plays an important role in IPv6 features such as

- Neighbor discovery (ND)
- · Path MTU discovery
- Multicast Listener Discovery (MLD)

You can identify an IPv6 ICMP packet because the Next Header field in the IPv6 packet header is 58

Internet Protocol Security (IPsec) with ICMPv6

You can configure IPsec with ICMPv6. For a configuration example of IPsec with ICMPv6, see Configuring Security for VOSS.

Path MTU discovery

IPv6 routers do not fragment packets.

The source node may send packets less than or equal to the maximum transmission unit (MTU) of the link layer.

As the packet travels through the network to the source it may encounter a link with a smaller MTU. If so, the router sends the source node an ICMP error message that contains the MTU size of the next link. The source IPv6 node then resends a packet equal to the size of the MTU included in the ICMP message.

The default Layer 3 IPv6 MTU value is 1500 where the system MTU default value is 1950.

The default IPv6 MTU value is always less than the default System MTU value.

You can configure the MTU for each IPv6 interface.

Note:

To configure separate Layer 3 MTU values for IPv4 and IPv6 packets on the same VLAN interface, you must disable Unicast Reverse Path Forwarding (uRPF) mode. If you enable the uRPF mode using the command boot config flags urpf-mode, the MTU values for both IPv4 and IPv6 packets on the same VLAN are matched. Different Layer 3 MTU sizes on the same VLAN are not allowed in uRPF mode.

Routing

A routing table is present on all nodes.

The routing table stores information about IPv6 network prefixes and how to reach them.



The switch requires routing protocols, such as OSPFv3 to exchange IPv6 routing prefixes.

For each incoming packet, the switch checks the destination neighbor cache first. If the destination is not in the destination neighbor cache, the routing table determines:

- the next-hop interface (the interface used for forwarding)
- · the next-hop address

Note:

The system uses the IPv6 Neighbor Cache for on-link, directly-connected destinations only. Off-link destinations go through a next-hop router, as determined by the next-hop address lookup.

IPv6 routes in a routing table can be:

- directly attached network routes using a 64-bit prefix
- remote network routes using a 64-bit or lower prefix
- · host routes using a 128-bit prefix length
- the default route using a prefix of ::/0

The switch supports the following IPv6 routing protocols:

- RIPng
- OSPFv3
- BGP+ (over 6in4 tunnels)
- IPv6 Shortcuts (over Fabric Connect)

You can redistribute IPv6 routes between any of these routing protocols.

This document focuses on OSPFv3. For information about OSPFv2, see Configuring OSPF and RIP for VOSS.

To configure IPv6 routing on a VLAN, an IP address is assigned to the VLAN. This IP address is not associated with any particular physical port, but is used on all ports where this VLAN is a member.

On a brouter port, a single port VLAN is used to route the traffic. IPv4 and IPv6 traffic is routed in the single-port brouter VLAN.

Other VLANs (which are multiple port VLANs) can bridge and route the traffic.

Virtual routing between IPv6 subnets

The switch supports IPv6 routing between subnets.

When you add an IP address to the VLAN, the system maps an IP subnet to the VLAN.

As shown in the following figure, although VLAN 1 and VLAN 2 reside on the same switch, for traffic to flow from VLAN 1 to VLAN 2, you must route the traffic.

You must enable IPv6 forwarding to route IPv6 traffic between VLANs. And you must enable IPv6 both globally and on a specific VLAN basis in order for forwarding to function. You can enable or disable IPv6, either globally or on a specific VLAN basis.

IPv6 forwarding is enabled by default.

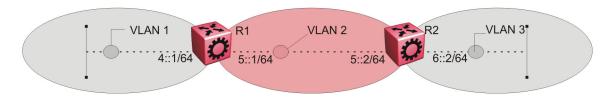


Figure 8: IPv6 routing between VLANs

When you configure routing on a VLAN, an IPv6 address assigned to the VLAN is the VLAN IP interface.

The VLAN IPv6 address can be reached through any VLAN port, and frames route from the VLAN through the gateway IPv6 address.

You can forward traffic to any IPv6 subnet in the switch. A VLAN can be reached only if it has an IPv6 interface configured on it.

Because a port can belong to multiple VLANs, a one-to-one correspondence no longer exists between the physical port and the router interface when VLAN tagging is enabled.

If you do not enable VLAN tagging a single port can belong only to one port-based VLAN, but that same single port can belong to multiple policy-based VLANs.

As with any IPv6 address, you can use any VLAN IP interface for device management.

For the Simple Network Management Protocol (SNMP) or Telnet management, you can use any VLAN IP interface to access the switch while routing is enabled on the VLAN.

Brouter ports

A brouter port is a single-port VLAN that can route IP packets and bridge all nonroutable traffic.

The difference between a brouter port and a standard protocol-based VLAN configured for routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic while it routes IP traffic.



Note:

Because a brouter port is a one-port VLAN, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

Static routes

Static routes provide an alternative method for establishing route reachability.

Static routes, with dynamic routes, provide routing information from the forwarding database.

Only enabled static routes whose nexthop address is reachable are submitted to the Route Table Manager (RTM), which determines the best route based on reachability, route preference, and cost.

The RTM communicates all updates to best routes.

If the nexthop is not reachable you can use the **show ipv6 route static** command to display the status. If the nexthop is not reachable, the status is <code>TryToResolve</code> and the route does not appear in the RTM until the nexthop address is resolved.

For directly-connected IPv6 Subnets you do not need to specify a nexthop address; you can specify outgoing Tunnel-ID, VLAN, or port. If you use outgoing Tunnel-ID, VLAN, or port, the implied nexthop value is 0::0.

When you configure IPv6 static routes only by interface (VLAN or brouter), it lets the traffic to reach IPv6 prefixes configured on the link that is directly connected to the interface provided in the static route configuration. For example: ipv6 route 180:0:0:0:0:0:0:0:0:0:0:0 1 vlan 631.

When you configure static routes with a link-local nexthop, you must also specify the outgoing Tunnel-ID, VLAN, or port because link-local addresses are ambiguous unless the proper interface binding is attached. For example: ipv6 route 1234::/64 cost 1 next-hop fe80::1 vlan 1900.

You must provide the following options to configure a static route:

local or nonlocal hop option

Configure a static route either with a next hop that exists on a locally attached network or a next hop that is reachable through a dynamic route. The static route is available as long as the next hop is reachable.

· route preference

You can specify the route preference for the static routes as follows:

- Global value for all static routes: the preference is either static or dynamic routes.
- Preference for each static route entry: if specified, this value overrides the global value for the entry which provides flexibility to change the general behavior of a specific static route.
- · Administrative status

Controls when the static route is considered for forwarding. Administrative status differs from the operational status. An admin-enabled static route can still be unreachable and not used for forwarding. An admin-disabled static route is operationally a nonexistent route.

· Multiple static routes

Specify alternative paths to the same destination. Multiple static routes provide stability and load balancing.

To configure a default static route, supply a value of 0 for the prefix and the prefix length.

The following table describes events that affect static route operation.

Table 7: Events and their affects on static route operation

Action	Result
Change the administrative status of the static route	Makes the static route unavailable for forwarding

Table continues...

Action	Result
	You can use one CLI command to administratively enable or disable all static routes as follows ipv6 route static enable.
	You can administratively disable all routes but preserve the static route configuration when you use the CLI command: no ipv6 static route enable.
Delete the IPv6 addresses of a VLAN or brouter port	Permanently deletes the static routes with the corresponding local neighbors from the RTM, the forwarding database, and the configuration database
Delete a VLAN	Removes static routes with a local next-hop option from the configuration database. Static routes with a nonlocal next-hop option become inactive (they are removed from the forwarding database).
Disable forwarding on a VLAN or brouter port	Static routes reachable through the locally attached network become inactive
Disable a VLAN or brouter port	Makes the static route inactive
Disable IPv6 forwarding globally	Stops forwarding all IPv6 traffic
Learn changes about a dynamically learned neighbor	After a neighbor becomes unreachable or is deleted, the static route with the neighbor becomes inactive, and the configuration is not affected. The static route with the neighbor becomes active in the configuration and is added to the RTM and forwarding database when the neighbor becomes reachable.
Enable a static route	Adds the route to the RTM to change certain static routes to active.
Delete a static route	Permanently deletes a static route from the configuration.
Disable a static route	Stops traffic on the static route but does not remove the route from the configuration.
Change a route preference	After the static route preference changes, the best routes for the entries use both static and dynamic paths.
Delete or disable a tunnel	Removes the tunnel entry from the forwarding table
Enable a tunnel	Activates the tunnel static routes and adds an entry to the forwarding table.

The local-nexthop flag is not required for IPv6.

An IPv4 device cannot learn a neighbor ARP entry unless the device uses a local route entry. In IPv6, a host can learn a neighbor entry if the device is physically connected to the neighbor (one hop).

The static route becomes active when the next hop is reachable by a dynamic route neighbor resolution. The static route takes the forwarding information from the dynamic route. If the next hop is reachable using a local route, the neighbor resolution is required.

Static route table

The static route table is separate from the system routing table that the router uses to make forwarding decisions.

You can use the static route table to directly change static routes.

Although the tables are separate, the system routing table automatically reflects the static routing table manager entries if the next-hop address in the static route is reachable and if the static route is enabled.

The static route table is indexed by four attributes:

- destination network
- destination mask
- next hop
- · interface

You can insert static routes by using the static route table, and you can delete static routes by using either the static route table or the system routing table. For information on route scaling, see Release Notes for VOSS.

Important:

The system routing table stores only active static routes with the best route preference. A static route is active only if the route is enabled and if the next-hop address is reachable; for example, if a valid IPv6 neighbor cache entry exists for the next hop.

You can enter multiple routes (for example, multiple default routes) that use different costs and the lowest cost route that is reachable appears in the routing table.

If you enter multiple next hops for the same route with the same cost, the switch does not replace the existing route.

If you enter the same route with the same cost and a different next hop, the switch uses the first route. If that first route becomes unreachable, the system activates the second route, with a different next-hop, with no connectivity loss.

Route scaling

IPv4 and IPv6 route scaling depends on the combination of the ipv6-mode and urpf-mode boot config flags. For more information, see Release Notes for VOSS.

IPv6 Circuitless IP

IPv6 Circuitless IP (CLIP) is a virtual interface that is not associated with any physical port. You can use an IPv6 CLIP interface to provide uninterrupted connectivity to your switch as long as an actual path exists to reach the device. The system treats the IPv6 CLIP interface like an IPv6 interface and treats the network associated with the IPv6 CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

You can use an IPv6 CLIP address as a logical IPv6 address for network management, as well as for other purposes. The IPv6 CLIP is typically a host address with any prefix length. You can redistribute this address as part of any other routing protocol update, so that the CLIP address is known to neighbors and available for use in routing or other types of connectivity. You can use IPv6 CLIP for many kinds of management connectivity such as telnet or SSH. You can also use IPv6 CLIP as a source IP address for sending Syslog messages.

For scaling information on IPv6 CLIP, see Release Notes for VOSS.

IPv6 CLIP restrictions and limitations

This section describes the restrictions and limitations associated with IPv6 CLIP.

- Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
- IPv6 CLIP does not support link-local address configuration.
- To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the IPv6 mode flag.



Note:

This flag does not apply to all hardware platforms. For more information, see Release Notes for VOSS.

- Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect duplicate IPv6 address assignment to this interface.
- Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
- IPv6 CLIP interface is enabled by default and it cannot be disabled.
- · You cannot configure an IPv6 CLIP interface as the source or destination endpoint of an IPv6in-IPv4 tunnel.

Equal Cost Multipath

Table 8: Equal Cost Multiple Path for IPv6 product support

Feature	Product	Release introduced	
For configuration details, see the following documents:			
Configuring IPv4 Routing for VOS	Configuring IPv4 Routing for VOSS		
Configuring IPv6 Routing for VOS	Configuring IPv6 Routing for VOSS		
Configuring BGP Services for VOSS			
ECMP for IPv6	VSP 4450 Series	VOSS 5.1	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 5.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 5.1	
	VSP 8400 Series	VOSS 5.1	
	VSP 8600 Series	VSP 8600 6.2	
	XA1400 Series	Not Supported	

With Equal Cost Multipath (ECMP), the switch can determine equal-cost paths to the same destination prefix. You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP traffic. ECMP is formed using routes from the same source or protocol, and the ECMP routes are evaluated within each routing protocol.

The ECMP feature supports and complements the following protocols and route types:

- OSPFv3
- · Static routes
- BGP+
- RIPng
- IPv6 Shortcuts

ECMP with static routes

ECMP supports and complements static routes.

The following points need to be considered while configuring ECMP with static routes:

 When ECMP is globally enabled, the equal cost static routes are added in the Routing Table Manager (RTM).

- Static routes that are configured only using an interface, such as VLAN or port, do not support ECMP as these routes have a preference value of 0 and are treated like local routes.
- If your switch supports a management interface, then static routes configured on the management interface do not support ECMP.
- Static routes configured by next-hop are not considered equal cost with routes that are configured by tunnel even if the routes have the same cost and preference.
- A static route configured by tunnel is the least preferred and is programmed only when a next-hop or an interface does not configure a static route.
- Static routes configured by next-hop that resolves their next-hop using another static route will be in the notReachable state even if the next-hop can be pinged.
- If there are two static routes configured by next-hop, and both next-hops are resolved via a dynamic protocol to the same value, then only one route will be in the reachableInRtm state. The state of the other route will be reachableNotInRtm.

Disable IPv6 ICMP multicast

On IPv6 networks, a packet can be directed to an individual machine or multicasted to a set of hosts. When a packet is sent to an IPv6 multicast address from a machine on the local network, that packet is delivered to a subset or all machines on that network.

If the packet that is sent to the multicast address is an ICMP Echo Request packet, the machines on the network will receive this ICMPv6 echo request packet and send an ICMP echo reply packet back. When all the machines on a network respond to this ICMPv6 echo request, the result can be severe network congestion or outages.

Network devices always respond to the IPv6 ICMP packets sent to a multicast address. However, you can disable the processing of IPv6 ICMP packets sent to a multicast address on the device. On disabling the ICMP multicast processing, all the packets containing ICMP sent to multicast addresses are dropped when the packets reach the control plane.

Route Policy Definition

You can define an IP route policy and its attributes globally, and then apply them individually to interfaces and protocols. You can also form a unified database of route policies that the RIP or OSPF protocol can use for type of filtering purpose. A name or ID identifies a policy.

Under a policy you can have several sequence numbers. If you do not configure a field in a policy, the field appears as 0 in CLI show command output. This value indicates that the device ignores the field in the match criteria. Use the clear option to remove existing configurations for the field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference

field, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce or redistribute purposes.

You can only apply one policy for each purpose (RIP Announce, for example) on a given RIP interface. In this case, all sequence numbers under the policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

The following tables display the accept, announce, and redistribute policies for RIP, OSPF, IS-IS and BGP. The tables also display which matching criteria apply for a certain routing policy. In these tables, 1 denotes advertise router, 2 denotes RIP gateway, and 3 denotes that external type 1 and external type 2 are the only options.

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.

Note:

IPv4 and IPv6 route-maps cannot be configured on the same match statement.

Table 9: Protocol route policy table for RIP

		Announce				
	OSPF	Direct	RIP	BGP	RIP	
Match Protocol	Yes	Yes	Yes	Yes		
Match Network	Yes	Yes	Yes	Yes	Yes	
Match IpRoute Source	Yes ¹		Yes ²			
Match NextHop	Yes	Yes	Yes	Yes	Yes	
Match Interface			Yes			
Match Route Type	Yes					
Match Metric	Yes	Yes	Yes	Yes	Yes	
MatchAs Path						
Match Community						
Match Community Exact						
MatchTag				Yes		
NssaPbit						

		Announce			
	OSPF	Direct	RIP	BGP	RIP
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type					
SetNextHop					
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					Yes
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 10: Protocol route policy table for OSPF

		Redistribute				
	Direct	Static	RIP	BGP	IS-IS	OSPF
Match Protocol				Yes	Yes	
Match Network	Yes	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes ²			
Match NextHop		Yes	Yes	Yes		
Match Interface			Yes			
Match Route Type					Yes ³	
Match Metric	Yes	Yes	Yes	Yes	Yes	Yes
MatchAs Path						
Match Community						
Match Community Exact						
MatchTag				Yes		
Set NSSA Bit	Yes	Yes	Yes	Yes	Yes	
SetRoute Preference						

		Redistribute				Accept
	Direct	Static	RIP	BGP	IS-IS	OSPF
SetMetric TypeInternal						
SetMetric	Yes	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes	
SetNextHop				Yes		
Set Inject NetList	Yes	Yes	Yes	Yes	Yes	Yes
SetMask						
SetAsPath						
SetAsPath Mode						
Set Automatic Tag						
Set CommunityNumber						
Set CommunityMode						
SetOrigin						
SetLocal Pref						
SetOrigin EgpAs						
SetTag						
SetWeight						

Table 11: Protocol route policy table for IS-IS

		Accept			
	Direct	Static	RIP	BGP	OSPF
Match Protocol	Yes	Yes	Yes	Yes	Yes
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes		
Match NextHop		Yes	Yes	Yes	Yes
Match Interface			Yes		
Match Route Type					Yes ³
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
Set NSSA Bit					
SetRoute Preference					
SetMetric Type Internal					

		Accept			
	Direct	Static	RIP	BGP	OSPF
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes
SetNextHop				Yes	
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 12: Protocol route policy table for BGP

		Redistribute			Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Match as-path				Yes	Yes
Match community	Yes	Yes	Yes	Yes	Yes
Match community-exact				Yes	Yes
Match extcommunity				Yes	Yes
Match interface					
Match local-preference					
Match metric	Yes	Yes	Yes	Yes	Yes
Match network	Yes	Yes	Yes	Yes	Yes
Match next-hop		Yes	Yes	Yes	Yes
Match protocol					
Match route-source				Yes	
Match route-type			Yes		Yes
Match tag					
Match vrf					
Match vrfids					

		Redistribute			Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Set as-path				Yes	Yes
Set as-path-mode				Yes	Yes
Set automatic-tag					
Set community				Yes	Yes
Set community-mode				Yes	Yes
Set injectlist	Yes	Yes	Yes		
Set ip-preference					
Set local-preference				Yes	Yes
Set mask					
Set metric	Yes	Yes	Yes	Yes	Yes
Set metric-type					
Set metric-type-internal					
Set next-hop				Yes	Yes
Set nssa-pbit					
Set origin					Yes
Set origin-egp-as					
Set Tag					
Set Weight				Yes	

IPv6 Basic Configuration using CLI

Use the procedures in this section to configure IPv6 basics using CLI.

Enabling the IPv6-mode boot config flag

Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits.

This flag is disabled by default. Use this procedure to enable (set to true) the IPv6-mode boot config flag.

When the IPv6-mode boot config flag is enabled, the maximum number of IPv4 routing table entries decreases. For scaling information, see Release Notes for VOSS.

Procedure

1. Enter Global Configuration mode:

enable

```
configure terminal
```

2. Enable the IPv6-mode boot config flag:

```
boot config flags ipv6-mode
```

- 3. Save the configuration, and then reboot the switch for the change to the IPv6-mode boot config flag to take effect.
- 4. After you reboot the switch, verify that the IPv6-mode boot config flag is set to true:

```
show boot config flags
```

Configuring an IPv6 static neighbor address

You can use static IPv6 neighbors to manually specify the link-layer address for a given IPv6 endpoint.

About this task

Under normal operation you do not need to configure static IPv6 neighbors.

However, IPv6 static neighbors can be used to:

- avoid the overhead associated with dynamic neighbor discovery protocol traffic
- help troubleshoot specific network scenarios

Note:

- IPv6 static neighbors are not supported on SMLT.
- When you add or remove IPv6 static neighbors that point to a nexthop router, make sure that you have disabled the IPv6 interface.

Not all parameters are available in non-default VRFs.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

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

2. Configure an IPv6 neighbor address:

```
ipv6 neighbor WORD < 0-128 > port {slot/port[sub-port]} mac < 0x00:0x00:0x00:0x00:0x00 > [vlan < 1-4059 > ]
```

- 3. Configure optional parameters if the default values do not meet your requirements:
 - a. Configure the hop limit:

```
ipv6 hop-limit <0-255>
```

The default is 64.

b. Configure ICMP network address unreachable messages:

```
ipv6 icmp addr-unreach
```

c. Configure the ICMP error interval:

```
ipv6 icmp error-interval <1-2147483647>
```

The interval is in milliseconds. An interval of 0 results in no error messages. The default is 1000.

d. Configure the ICMP error quota:

```
ipv6 icmp error-quota <0-2000000>
```

The default is 50.

e. Configure ICMP port unreachable messages:

ipv6 icmp port-unreach

f. Enable response to icmp echo multicast packets:

```
ipv6 icmp echo-multicast-request
```

The default is disabled.

g. Enable ICMP network unreachable messages:

```
ipv6 icmp unreach-msg
```

The default is disabled.

Example

Add an IPv6 neighbor for a brouter port:

```
Switch:1(config) #ipv6 neighbor 3000:0:0:0:0:0:0:2 port 1/11 mac 00:0c: 42:07:35:90
```

Add an IPv6 neighbor for a VLAN:

Switch:1(config) #ipv6 neighbor 3000::3 port 1/12 mac 01:02:03:04:05:06 vlan 20

Variable definitions

Use the data in the following table to use the ipv6 commands in this procedure.

Variable	Value
forwarding	Configures whether this entity is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, this entity. Enable forwarding to act as a router. The default is enabled.
hop-limit <0–255>	Configures the hop limit. The default is 64.

Variable	Value			
icmp addr-unreach	Enables ICMP address unreachable messages.			
	The default is enabled.			
icmp echo-multicast-request	Enable response to icmp echo multicast packets.			
	The default is enable.			
icmp error interval <1-2147483647>	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000.			
icmp error-quota <0-2000000>	Configures the number of ICMP error messages that can be sent during the ICMP error interval.			
	A value of zero instructs the system not to send an ICMP error messages.			
	The default value is 50.			
icmp port-unreach	Enables ICMP port unreachable messages.			
	The default is enabled.			
icmp unreach-msg	Enables ICMP network unreachable messages. The default is disabled.			
neighbor WORD<0-128> port {slot/port[sub-port]} mac <0x00:0x00:0x00:0x00:0x00:0x00> [vlan	Creates a static IPv6 neighbor with the following variables:			
<1-4059>]	WORD<0-128> specifies the IPv6 addressof the neighbor in hexadecimal colon format.			
	• {slot/port[/sub-port]} specifies the brouter port to use for the neighbor.			
	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.			
	mac <0x00:0x00:0x00:0x00:0x00:0x00> specifies the MAC address of the neighbor.			
	• vlan <1-4059> specifies the VLAN ID to use for the neighbor.			
	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.			

Variable	Value
	Static IPv6 neighbors do not maintain any state machine and the system assumes that they are always reachable.

Configuring an IPv6 interface

The information in this section can help you configure an IPv6 interface to make IPv6 active on the interface and fine-tune IPv6 neighbor discovery to control the frequency of protocol traffic.

By default, IPv6 forwarding is enabled on an interface.

Compared to IPv4/ARP, the IPv6 neighbor discovery mechanism maintains more protocol state, timers, and protocol traffic overhead.

There are two important tunable parameters for IPv6 ND that can control the frequency of protocol traffic:

- ipv6 interface reachable-time
- ipv6 interface retransmit-timer

Before you begin

 Before you can assign an IPv6 address to the interface, you must configure an IPv6 interface for a VLAN or brouter port.

You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address.

The switch supports port-based and IPv6 protocol-based VLANs.

For information about how to configure VLANs, see the following documents:

- Configuring VLANs, Spanning Tree, and NLB for VOSS
- Configuring Link Aggregation, MLT, SMLT and vIST for VOSS

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Create IPv6 interface:

```
ipv6 interface
```

- 3. Configure optional parameters to meet your requirements:
 - a. Enable IPv6 router advertisement on the interface:

```
ipv6 nd send-ra
```

b. Configure the maximum number of hops before packets drop:

```
ipv6 interface hop-limit <1-255>
```

c. Configure the link-local address:

```
ipv6 interface link-local WORD<0-19>
```

d. Configure the mac offset:

```
ipv6 interface mac-offset <MAC-offset>
```

e. Configure the maximum transmission unit (MTU):

```
ipv6 interface mtu <1280-9500>
```

f. Configure an interface description:

```
ipv6 interface name WORD<0-255>
```

g. Configure the time a neighbor is considered reachable after receiving a reachability confirmation:

```
ipv6 interface reachable-time <1-3600000>
```

h. Configure the time between retransmissions of neighbor solicitation messages to a neighbor when resolving the address or when probing the reachability of a neighbor:

```
ipv6 interface retransmit-timer <1-4294967295>
```

i. Configure a brouter port as part of an IPv6 VLAN:

```
ipv6 interface vlan <1-4059>
```

- j. Configure the interface to perform IPv6 unicast reverse path forwarding:
 - To enable urpf-mode boot flag, enter:

```
boot config flags urpf-mode
```

To configure unicast reverse path forwarding, enter:

```
ipv6 rvs-path-chk mode {strict|exist-only}
```

Example

Create and administratively enable the interface:

```
Switch:1(config-if) #ipv6 interface enable
```



Note:

In contrast to IPv6 interface creation and address assignment in EDM, you use the ipv6 interface CLI command to create an interface and specify a single global address in one step.

Variable definitions

Use the data in the following table to use the ipv6 interface command.

Variable	Value
hop-limit <1–255>	Configures the maximum hops. The default is 64.
link-local WORD<0-19>	Specifies the 64-bit interface ID used to calculate the actual link-local address as a name up to 19 characters long.
mac-offset <mac-offset></mac-offset>	Use mac-offset to request a particular MAC for an IPv6 VLAN.
	Note:
	This parameter applies only to VLANs.
	You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated. Different hardware platforms support different ranges. To see which range is available on the switch, use the CLI command completion Help.
mtu <1280–9500>	Configures the maximum transmission unit for the interface: 1280–1500, 1850, or 9500. This value must be the same for all addresses defined on this interface.
	The default is 1500.
	Different hardware platforms support different MTU values. To see what values your switch supports, use the CLI command completion help.
name WORD<0-255>	Assigns a descriptive name. The network management system also configures this string.
reachable-time <0-3600000>	Controls how long IPv6 neighbor entries learned on an interface remain in the REACHABLE state (as described in RFC 4861).
	The system randomizes the value you configure, per RFC specifications, to be 50%-150% of the configured value.
	By default the reachable-time base value is 30 seconds, with an actual 15-45 second range when you consider the randomization factor.

Variable	Value
	The default is 3000 milliseconds
retransmit-timer <0-4294967295>	Controls the time, in milliseconds, between retransmission of Neighbor Solicitation messages when the system attempts to resolve or reconfirm the reachability of an IPv6 neighbor.
	By default, the system sends three Neighbor Solicitation messages with a one second interval between each message. If the system does not receive a corresponding Neighbor Advertisement within an interval equal to 3 X retransmit-timer milliseconds, the system declares the IPv6 neighbor unreachable.
	Tip:
	You can increase the retransmit-timer to extend the interval that the switch waits until it declares the neighbor unreachable. For example: a retransmit-timer value of 5000 means the switch waits 3 X 5000 milliseconds which equals 15000 milliseconds or 15 seconds.
	The default is 1000 milliseconds
vlan <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 value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.

Use the data in the following table to use the ipv6 rvs-path-chk command.

Variable	Value
mode {strict exist-only}	Specifies the mode for Unicast Reverse Path Forwarding (uRPF).
	In strict mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. If the incoming interface is not the best reverse path, the packet check fails and uRPF drops the packet.
	In exist-only mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. The packet is dropped only if the source address is not reachable via any interface on that router.

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

Assigning IPv6 addresses to a brouter port or VLAN

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Assign an IPv6 address:

```
ipv6 interface address WORD<0-255>
```

Example

Assign an IPv6 address specifying the full 128 bits of the address:

```
Switch:1(config-if) #ipv6 interface address 30:0:0:0:0:0:0:0:fffff/64
```

Assign an IPv6 address specifying only the upper 64 bits of the address:

```
Switch:1(config-if)#ipv6 nd prefix-interface <prefix> eui <1-3>
```

In this example you specify only the upper 64 bits of the address and allow the system to autogenerate the lower 64 bits from the MAC address.

Variable definitions

Use the data in the following table to use the ipv6 interface address command.

Variable	Value
WORD<0-255>	Specifies the IPv6 address for the port or VLAN.

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring IPv6 route preferences

Before you begin



! Important:

Changing route preferences can affect system performance and network accessibility while you perform the procedure. Change a prefix list or a routing protocol before you activate the protocols.

About this task

Configure IPv6 route preferences to give preference to routes learned for a specific protocol.

Procedure

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

```
enable
configure terminal
```

Optional: router vrf WORD<1-16>

2. Configure the route preference:

ipv6 route preference protocol <static|ebgp|ibgp|ospfv3-intra|
ospfv3-inter|ospfv3-extern1|ospfv3-extern2|ripng|spbm-level1>
<0-255>

3. Confirm that the configuration is correct:

show ipv6 route preference [vrf WORD<1-16> | vrfids WORD<0-512>]

Example

Configure the route preference to RIPng and confirm the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ipv6 route preference protocol ripng 10
Switch:1(config)#show ipv6 route preference

----

IPv6 Route Preference - GlobalRouter

----

PROTOCOL DEFAULT CONFIG

----

LOCAL 0 0 0
STATIC 5 5
SPBM_L1 7 7
OSPFv3_INTRA 20 20
OSPFv3_INTRA 20 20
OSPFv3_INTER 25 25
EBGP 45 45
RIPNG 100 100
OSPFv3_E1 120 120
OSPFv3_E2 125 125
IBGP 175 175
```

Variable definitions

Use the data in the following table to use the ipv6 route preference and the show ipv6 route preference commands.

Variable	Value
<0-255>	Assigns a route preference value.
ebgp	Configures the preference for protocol type EBGP.
ibgp	Configures the preference for protocol type IBGP.
ospfv3-extern1	Configures the preference for protocol type OSPFv3 external type 1.
ospfv3-extern2	Configures the preference for protocol type OSPFv3 external type 2.
ospfv3-intra	Configures the preference for protocol type OSPFv3 intra-area.
ospfv3-inter	Configures the preference for protocol type OSPFv3 inter-area.
ripng	Configures the preference for protocol type RIPng.

Variable	Value
spbm-level1	Configures the preference for protocol type spbm-level1.
static	Configures the preference for protocol type static.

Viewing global IPv6 information

View and manage general IPv6 information.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display IPv6 information for an interface:

```
show ipv6 interface [gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}] [loopback <1-256>][mgmtEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}][tunnel <1-2000>][vlan <1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

3. Display IPv6 tunnel information:

```
show ipv6 interface tunnel <1-2000>
```

4. Display IPv6 address information for the specified slot and port:

```
show ipv6 address interface gigabitethernet {slot/port[/sub-port] [-
slot/port[/sub-port]] [,...]} [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

5. Display IPv6 address information for the specified IPv6 address:

```
show ipv6 address interface ip WORD < 0-46 > [vrf WORD < 1-16 > ] [vrfids WORD < 0-512 > ]
```

6. Display IPv6 address information for the specified tunnel:

```
show ipv6 address interface tunnel <1-2000>
```

7. Display IPv6 address information for the specified VLAN:

```
show ipv6 address interface vlan <1-4059>
```

8. Display the current state of IPv6 forwarding:

```
show ipv6 forwarding [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

9. Display information on the current state of IPv6 functionality:

```
show ipv6 global [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

10. Display IPv6 Gigabit Ethernet (GbE) router advertisement information:

```
show ipv6 nd interface gigabitethernet [{slot/port[/sub-port] [-
slot/port[/sub-port]] [,...]}]
```

11. Display IPv6 router advertisement information for the management port:

show ipv6 nd interface mgmtEthernet mgmt



This step only applies to hardware with a dedicated, physical management interface.

12. Display IPv6 VLAN router advertisement information:

```
show ipv6 nd interface vlan [<1-4059>]
```

13. Display detailed information in IPv6 router advertisements:

```
show ipv6 nd-prefix detail [vrf WORD<1-16> | vrfids WORD<0-512>]
```

14. Display GbE interface information in IPv6 router advertisements:

```
show ipv6 nd-prefix interface gigabitethernet [{slot/port[/sub-port]
[-slot/port[/sub-port]] [,...]}]
```

15. Display VLAN interface information in IPv6 router advertisements:

```
show ipv6 nd-prefix interface vlan [<1-4059>]
```

16. Display VLAN information in IPv6 router advertisements:

```
show ipv6 nd-prefix vlan <1-4059>
```

17. Display IPv6 neighbor entries with specific brouter port numbers:

```
show ipv6 neighbor interface gigabitethernet {slot/port[/sub-port]}
```

18. Display IPv6 neighbor information for neighbors of the specified type:

```
show ipv6 neighbor type <1-4>
```

19. Display IPv6 neighbor information:

show ipv6 neighbor [WORD<0-46>]

Example

Switch:	1 (cor	nfig)#show	ipv6 in	terface '	vlan										
======			======	======	=====	Vlan	Ipv6	Inte	erface	====	=======	======	=======		
IFINDX INDX		PHYSICAL ADDRESS		ADMIN STATE	OPER STATE	TYPE			REACHABI TIME	LΕ	RETRANSMIT TIME	MCAST STATUS	IPSEC	RPC	RPCMODE
2070	22	e4:5d:52:	3c:65:02	disable	down	ETHER	1500	64	30000		1000	disable	disable	disable	existonly
======	=====		======	Vlan I	====== pv6 Ad	dress	=====	====	======	====					
IPV6 AD	DRESS	3			VLA	VLAN-ID		PE.	ORIGIN	1	STATUS				
fe80:0:	0:0:6	e65d:52ff:	fe3c:650	2	V-2	2	UI	NICAS	ST LINKLA	AYER	INACCESSIBL	E			
		Total Num													
Switch:	1#shc	ow ipv6 in	terface	tunnel											
				Tun	nel Ip	v6 Int	erface	e - 0	lobalRοι	ıter					
IF De	escr	VL.	AN PHYSI ADDRE			DMIN TATE	OPER STATE				REACHABLE TIME	RETRANSM TIME	IT MCAST STATU		
6656 T-	-1	0	n/a		e	nable	down	P21	1280	64	30000	1000	disab	le	

```
Tunnel Ipv6 Address
                                                          TUNNEL-ID TYPE ORIGIN STATUS

T-1 UNICAST MANUAL INACCESSIBLE INF INF
T-1 UNICAST LINKLAYER INACCESSIBLE INF INF
IPV6 ADDRESS
11.0.0.0.0.0.0.11.11/32
fe80:0:0:0:0:0:b0b:b0b/64
1 out of 5 Total Num of Interface Entries displayed.
2 out of 9 Total Num of Address Entries displayed.
Switch: 1#show ipv6 address interface tunnel 2
______
                                                                      Address Information
_____
                                                                                              VID/BID/TID TYPE ORIGIN STATUS
IPV6 ADDRESS/PREFIX LENGTH
             ______
                                                                                             T-2 UNICAST MANUAL PREFERRED
T-2 UNICAST LINKLAYER PREFERRED
44:211:0:0:0:0:0:2/64
fe80:0:0:0:0:0:d301:3702/64
2 out of 407 Total Num of Address Entries displayed.
Switch: 1#show ipv6 address interface vlan 100
______
                                                                      Address Information
 IPV6 ADDRESS/PREFIX LENGTH
                                                                                             VID/BID/TID TYPE ORIGIN STATUS
                                                                                            V-100 UNICAST MANUAL PREFERRED
V-100 UNICAST LINKLAYED DEFERDED
10:1:50:0:0:0:0:7/64
fe80:0:0:0:b2ad:aaff:fe46:f19a/64
2 out of 407 Total Num of Address Entries displayed.
Switch:1#show ipv6 forwarding
            Ipv6 forwarding - GlobalRouter : enable
                                       : disable : 1
            ecmp
          ecmp-max-path
Switch:1#show ipv6 global
                             IPv6 Global Information - GlobalRouter
_____
                               : disable
      forwarding
     forwaruing default-hop-cnt : 64 number-of-interfaces : 1 icmp-error-interval : 1000 icmp-error-quota : 50
     | 1cmp-error-interval | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 | 1000 |
      ecmp-max-path
    source-route : disable host-autoconfig : disable
Switch:1#show ipv6 nd interface vlan
                                                        Vlan Ipv6 Nd - GlobalRouter
IFID VLAN RTR-ADV MAX-INT MIN-INT LIFETIME MANAG OTHER DAD-NS MTU HOP
REACHABLE RETRANSMIT
           TIME
TIME
2092 V-44 True 600 200 1800 False False 1 0(d) 64(d) 30000(i) 1000(i) 2081 V-33 True 600 200 1800 False False 1 0(d) 64(d) 30000(i) 1000(i)
Note: (s) = Set, (d) = Default, (i) = inherit.
2 out of 2 Total Num of Ipv6 ND Entries displayed.
Switch:1#show ipv6 nd interface mgmtEthernet mgmt
                                                                    Mgmt Ipv6 Nd
```

IFID MGMT-	-IF RTR-AD	OV MAX-I	NT MIN-IN	T LIFETIME	E MANA	G OTHER CONF	DAD-NS M	TU H LI	OP F	REACHABI	LE RETRANSMIT TIME	
				1800								
Switch:1#s	show ipv6	nd-pref	ix interf	ace gigabi	itethe	rnet						
=======				Ipv6 Nd Pr								
INTF IPV6												
INDEX ADDR												
344 2011	:beef:400	04:0:0:0	0:0:0/64				5/25 25	92000	604800) 1		
1 out of 9	Total Nu	um of Ip	ov6 ND pre	efix Entrie	es disp	played.						
Switch:1#s	show ipv6	nd-pref	ix interf	ace vlan								
				Nd D						-===		
				Ipv6 Nd Pr 						-===		
INTF IPV6		IX					VALID LIFE					
2148 2011			0:0/64				2592000 2592000					
2158 2011 2248 2011	:beef:200	0:0:0:0:	0:0/64			200	2592000	604	800 1			
2258 2011 2548 2011	:beef:500	0:0:0:0:	0:0/64			500	2592000 2592000 2592000 2592000 2592000	604	800 1			
2648 2011 2948 2011						900	2592000 2592000	604 604	800 1			
7 out of 9	Total Nu	um of Ip	v6 ND pre	efix Entrie	es disp							
Switch:1#s	show ipv6	nd-pref	ix vlan 1	.00								
				Address Ir								
INTF IPV6							VALID					
INDEX ADDR	RESS/PREFI					ID	LIFE	LIFE				
2148 2011	:beef:100	0:0:0:0:	0:0/64			100	2592000	604800	1			
Legend: EU										compleme	ent(3)	
Switch:1#s	show ipv6	neighbo	or type 4									
				Neighbor	Info	rmation	- Global	Router				
NET ADDRES	SS/ ADDRESS			IPV6 INTE	5 PH:	YS FF		TYPE	STATE	L) Ul		
123:0:0:0: 00:00:	0:0:0:123	3/ :01			срі	P		LOCAL				
1 out of 1												
Switch:1(c			_	_	rayed	•						
				Neighbor	r Info	rmation	- Global	Router				
NET ADDRES	SS/ ADDRESS			IPV6 INTE	5 PH' F IN'	YS FF		TYPE	STATE	L) Ul	AST TUNNEL PD	
22:0:0:0:0	0:0:0:22/ 71:49:38:				2 1/2	2		DYNAMIC LOCAL	REACHA	ABLE 5	70	
	0e:c8:8c:		2002/					DYNAMIC				
84:83: fe80:0:0:0	71:49:38: cd6ea:eff	:82 :fec8:8		V-22 V-22				LOCAL				
fe80:0:0:0		:fec8:8									23	

Variable definitions

Use the data in the following table to use the **show ipv6** commands.

Variable	Value
address interface ip WORD<0-46>	Specifies the IPv6 address.
neighbor [WORD<0-46>]	Specifies the IPv6 address of the neighbor.
loopback <1-256>	Specifies the loopback interface ID value. If you do not specify a value, the output includes all IPv6 loopback interfaces.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
type <1-4>	Specifies the neighbor type as one of the following:
	• 1 - other
	• 2 - dynamic
	• 3 - static
	• 4 - local
tunnel <1-2000>	Specifies the tunnel ID.
vlan <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.
vrfWORD<1-16>	Specifies the VRF name.
vrfidsWORD<1-256>	Specifies the VRF ID.

Creating IPv6 static routes

Use static routes to manually configure routes to destination IPv6 address prefixes.

About this task

Not all parameters are available in non-default VRFs.

Procedure

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

```
enable
configure terminal
```

Optional: router vrf WORD<1-16>

2. Enable IPv6 static routes globally:

```
ipv6 route static enable
```

If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM.

3. Configure a static route:

```
ipv6 route WORD<0-46> [enable] [cost <1-65535>] [next-hop WORD<0-46>] [preference <1-255>] [tunnel <1-2000>] [port \{slot/port[subport]\}] [vlan <1-4059>]
```

4. (Optional) Disable all IPv6 static routes:

no ipv6 route static enable

5. **(Optional)** Permanently delete the IPv6 static route configuration:

clear ipv6 route static [vrf WORD<1-16> | vrfids WORD<0-512>]

Example

Enable IPv6 static routes globally:

Switch:1(config) #ipv6 route static enable

Create and enable a static route through a global nexthop:

Switch:1(config) #ipv6 route 4000::/64 cost 1 next-hop 3000::2 enable

Create and enable a static route through an outgoing interface (VLAN or brouter port):

Switch:1(config) # ipv6 route 4000::/64 cost 1 vlan 1900 enable

Create and enable a static route through a link local nexthop and an outgoing interface:

Switch:1(config) # ipv6 route 4000::/64 cost 1 next-hop fe80::1 vlan 1900 enable

In the preceding example, you must specify the outgoing interface so that the system can apply the correct context to the link-local address.

Variable definitions

Use the data in the following table to use the ipv6 route command.

Variable	Value
WORD <0-46>	Specifies the IPv6 destination address and prefix.
enable	Enables the static route. The default is enabled.
cost <1-65535>	Specifies the cost or distance ratio to reach the destination for this static route. The default is 1.

Variable	Value
next-hop Word <0-46>	Specifies the IPv6 address of the next hop on this route. You do not need to specify the next hop if the devices directly connect to one another. Configure the next hop if the two nodes do not share the same network prefix but reside on the same link.
preference <1–255>	Specifies the routing preference of the destination IPv6 address. The default is 5.
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
tunnel <1-2000>	Specifies the tunnel ID.
vlan <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.

View Routes Information

View routes information to view the current configuration.

About this task

Not all parameters are available in non-default VRFs.

IPv6 host routes created for the IPv6 local interfaces do not display in the routing table.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Show route information for alternative routes:

```
show ipv6 route alternative [vrf WORD<1-16> | vrfids WORD<1-256>]
```

3. Show the number of OSPF, RIP, static, and local routes:

```
show ipv6 route count-summary [vrf WORD<1-16> | vrfids WORD<1-256>]
```

4. Show route information for a destination:

```
show ipv6 route dest WORD < 0-46 > [vrf WORD < 1-16 > | vrfids WORD < 1-256 > ]
```

5. Show route information for a port:

```
show ipv6 route gigabitethernet {slot/port[sub-port]}
```

6. Show route information for a next-hop address:

show ipv6 route next-hop WORD<0-46> [vrf WORD<1-16> | vrfids WORD<1-256>]

7. Show route information for an SPBM IPv6 route:

show ipv6 route spbm-nh-as-mac

8. Show route information for a static route:

show ipv6 route static [vrf WORD<1-16> | vrfids WORD<1-256>]

9. Show route information for a tunnel:

show ipv6 route tunnel <1-2000>

10. Show route information for a VLAN:

show ipv6 route vlan <1-4059>

Example

	IPv6 Routing Table Informa	ation - GlobalRouter					
Destination Address/PrefixLen PREF	NEXT HOP	NH VRF/ISID	VID/BID/TID	PROTO	COST	AGE	TYPE
22:0:0:0:0:0:0:0/64	0:0:0:0:0:0:0:0	-	V-22	LOCAL	1	0	В
0 123:0:0:0:0:0:0:0/64 0	0:0:0:0:0:0:0:0	-	C-1	LOCAL	1	0	В
4 out of 4 Total Num of Route Entrie TYPE Legend: A=Alternative Route, B=Best Route, E							

Switch:1#show ipv6 route count-summary

IPv6 Route Summary - GlobalRouter

VRF NAME TOTAL OSPF RIP BGP STATIC LOCAL ISIS

GlobalRouter 13 10 0 3 1 2 7

Static Route Information - GlobalRouter

DEST-IP
NEXT-HOP
NEXT-HOP
PREFERENCE

22:0:0:0:0:0:0:66

1 out of 1 Total Num of Static Routes displayed.

Global IPv6 Static Routes Admin Status: enable

Switch:1#show ipv6 route spbm-nh-as-mac

IPv6 Routing Table Information - GlobalRouter

2001:cdab:0:0:0:0:0/32	Destination Address/PrefixLen	NEXT HOP	NH VRF/ISID	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2002:cdah:0.0.0.0.0.0/32 80.2d.30.00.00.01 - V-10 TSTS 1 0 B 0	2001:cdab:0:0:0:0:0:0/32	0:0:0:0:0:0:0:0	-	V-611	LOCAL	1	0	В	0
2002.000.010.010.010.01	2002:cdab:0:0:0:0:0:0/32	80:2d:30:00:00:01	-	V-10	ISIS	1	0	В	0

Variable definitions

Use the data in the following table to use the **show ipv6** route command.

Variable	Value
count-summary	Shows the total number of OSPF, static, and local routes.
dest WORD<0-46>	Specifies the IPv6 destination network address. The prefix value must match the prefix length.
next-hop WORD<0-46>	Specifies the IPv6 address of the next hop on this route.
spbm-nh-as-mac	Shows the B-MAC address as the next hop rather than the host name.
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
static	Shows static IPv6 routes.
tunnel <1-2000>	Shows route entries for a specific tunnel ID.
vlan<1-4059>	Shows route entries for a specific VLAN ID.
vrfWORD<1-16>	Shows the VRF name.
vrfidsWORD<1-256>	Shows the VRF ID.

Creating an IPv6 CLIP interface

About this task

Perform this procedure to create an IPv6 CLIP interface and associate it with a specific VRF.

Procedure

1. Enter Loopback Interface Configuration mode

```
enable
configure terminal
interface Loopback <1-256>
```

2. Create an IPv6 loopback interface address:

```
ipv6 interface address WORD <0-255> vrf WORD <1-16>
```

3. Ensure the configuration is correct:

show ipv6 interface loopback <1-256>

Example

				Loopback	Ipv6 Inte	erfaces							
IF INDX	VRF NAME	Descr		PHYSICAL ADDRESS	ADMIN STATE	OPER STATE	TYPE	MTU		REACHAB: TIME		======= RETRANSMIT FIME	IPSEC STATE
1348	GlobalRouter	CLIPv6-5		00:00:00:00:00:0	5 enable	up	ETHER	1500	64	30000	:	1000	disable
				oopback IPv6 Add	resses							===	
IPV6	ADDRESS/PREFIX	LENGTH		LOOPBACK-ID	TYPE	ORIGIN	ST	ATUS					
2001	:db8:0:0:0:0:0:f	fff/32		C-5	UNICAST	MANUAL	PRI	EFERRI	ED II	1F	INF		
Legei	nd: NA - Informa	tion not avail	lable										
	t of 1 Total Num t of 1 Total Num												

Variable definitions

Use the data in the following table to use the ipv6 commands.

Variable	Value
WORD<1-256>	Specifies the CLIP interface ID.
WORD<0-255>	Specifies the IPv6 address.
vrf WORD<1–16>	Specifies the VRF name.

Enabling IPv6 ECMP

About this task

Use the following procedure to enable IPv6 ECMP globally or on a specific VRF. IPv6 ECMP is disabled by default.

Procedure

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

enable
configure terminal

Optional: router vrf WORD<1-16>

2. Enter the following command to enable IPv6 ECMP:

ipv6 ecmp enable

3. Set the default value, IPv6 ECMP is disabled by default.

default ipv6 ecmp enable

4. Disable IPv6 ECMP globally:

no ipv6 ecmp enable

Variable definitions

Use the data in the following table to use the ipv6 ecmp command.

Variable	Value
enable	Enables IPv6 ECMP globally.
	Note:
	 Enabling IPv6 ECMP sets the maximum number of paths configured to its default value. This value is either 4 or 8 depending on your hardware platform.
	 Disabling IPv6 ECMP sets the maximum number of paths configured to 1.

Configuring maximum number of ECMP paths

Before you begin

Enable ECMP on the switch before configuring the max-path value. For more information, see ECMP on page 61.

About this task

Use the following procedure to configure the maximum number of ECMP paths.

Procedure

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

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

2. Configure the maximum number of ECMP paths.

```
ipv6 ecmp max-path <ECMP-Paths>
```

3. Set the configured maximum number of ECMP paths to its default value:

```
default ipv6 ecmp max-path
```



The default value for max-path is the maximum value, which varies depending on your hardware platform.

Example

Switch:1>enable Switch:1#configure terminal

Enable ECMP on the switch:

Switch:1(config) #ipv6 ecmp enable

Configure the maximum number of ECMP paths to 4:

Switch:1(config) #ipv6 ecmp max-path 4

Variable definitions

Use the data in the following table to use the ipv6 ecmp max-path command.

Variable	Value
<ecmp-paths></ecmp-paths>	Specifies the maximum number of ECMP paths. Different hardware platforms can support a different number of ECMP paths. For more information on the maximum number of ECMP paths supported on the switch, see the scaling information in Release Notes for VOSS.
	When ECMP is enabled, the default value is either 4 or 8 depending on your hardware platform.

Enabling or disabling IPv6 ICMP multicast

On disabling the ICMP multicast processing, ICMPv6 Echo Request packets sent to IPv6 multicast addresses are dropped when they reach the control plane.

About this task

Use this procedure to enable or disable the IPv6 ICMP multicast on the global router.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

enable

configure terminal

Optional: router vrf WORD<1-16>

2. Enable IPv6 ICMP multicast, enter:

ipv6 icmp echo-multicast-request

3. Disable IPv6 ICMP multicast, enter:

no ipv6 icmp echo-multicast-request

4. Set IPv6 ICMP multicast to default state, enter:

default ipv6 icmp echo-multicast-request



Note:

By default, the IPv6 ICMP multicast feature is enabled.

5. View the IPv6 ICMP multicast state:

```
show ipv6 global [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Enabling Stateless Address Autoconfiguration

Enable IPv6 Stateless Address Autoconfiguration to generate addresses using a combination of locally available information and information advertised by routers.

The default is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable IPv6 autoconfiguration:

```
ipv6 autoconfig
```

Example

```
Switch:1(config) #ipv6 autoconfig
```

The following example shows a sample output for the show ipv6 global command.

```
Switch: 1#show ipv6 global
                                    : disable
       forwarding
       default-hop-cnt
                                     : 64
       number-of-interfaces
                                    : 2
       icmp-error-interval
                                    : 1000
       icmp-error-quota
                                     : 50
       icmp-unreach-msg
                                     : disable
       icmp-addr-unreach-msg
icmp-port-unreach-msg
                                    : enable
                                    : enable
       icmp-echo-multicast-request : enable
       static-route-admin-status
                                     : enable
       alternative-route
                                     : enable
       ecmp
                                     : disable
       ecmp-max-path
                                     : 1
                                     : disable
       source-route
       host-autoconfig
                                    : enable
```

Configuring route advertisement on the management port

Configure route advertisement in IPv6 on the management port for neighbor discovery (ND).

Note:

Not all hardware platforms include a dedicated, physical management interface. For more information about supported interfaces, see your hardware documentation.

Procedure

1. Enter mgmtEthernet Interface Configuration mode:

```
enable
configure terminal
interface mgmtEthernet <mgmt | mgmt2>
```

2. Configure the number of neighbor solicitation messages from duplicate address detection:

```
ipv6 nd dad-ns <0-600>
```

3. Configure the hop limit sent in router advertisements:

```
ipv6 nd hop-limit <0-255>
```

Example

Variable definitions

Use the data in the following table to use the ipv6 nd command.

Variable	Value
dad-ns <0-600>	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD).
	A value of 0 disables the DAD process on this interface.
	A value of 1 sends one advertisement without retransmissions.
	The default is 1.
hop-limit <0–255>	Specifies the current hop limit field sent in router advertisements from this interface.
	The value must be the current diameter of the Internet.

Variable	Value
	A value of zero indicates that the advertisement does not specify a hop-limit value.
	The default is 64.

Configuring process-redirect for management port

Configure process-redirect messages to honor or ignore redirect messages for the management port. Redirect messages are visible only when Stateless Address Autoconfiguration is configured on switches capable of routing IPv6 traffic.

The default is disabled.



Note:

Not all hardware platforms include a dedicated, physical management interface. For more information about supported interfaces, see your hardware documentation.

Before you begin

- Disable IPv6 forwarding.
- Enable Stateless Address Autoconfiguration
- Create an IPv6 interface.
- · Configure an IPv6 address.

Procedure

1. Enter mgmtEthernet Interface Configuration mode:

```
enable
configure terminal
interface mgmtEthernet <mgmt | mgmt2>
```

2. Configure process-redirect messages:

```
ipv6 interface process-redirect
```

3. Verify that process-redirect messages are configured on the management port.

```
Switch:1(config-if) #ipv6 interface process-redirect
```

Example

The following examples shows a sample output of the show ipv6 interface processredirect command.

Switch:1#show ip	v6 interface process	-redirect	
	Process ICMP r	edirect status	
IFINDX	DESCR	VLAN	STATUS
64	PORT-mgmt	4092	Enabled
192	VLAN-5	5	Disabled

2050 VLAN-2 2 Disabled

Viewing IPv6 default routers

View the table of default routers learned from router advertisement messages.

A maximum of four routers are visible in the default routers list.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Show IPv6 default routers:

show ipv6 default-routers

Example

Switch:1#show ipv6 default-route	rs		
Default Route	rs VLAN	LIFETIME	IS ACTIVE
	mgmt mgmt	1778 1657	YES NO

Job aid

The following table describes the fields in the output for the **show ipv6 default-routers** command.

Parameter	Description
net address	Shows the IPv6 router address received from a valid router advertisement.
VLAN <1-4059>	Shows 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. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
LIFETIME	Shows the value placed in the router lifetime field of router advertisements. This value must be either 0 or between 4 and 9000. A value of zero indicates that the system is not a default router. The default is 1800.
IS ACTIVE	Shows if the default router is active or inactive.

Configuring an IPv6 prefix list

Use IPv6 prefix lists to allow or deny specific IPv6 route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

About this task

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create an IPv6 prefix list:

```
ipv6 prefix-list <WORD 1-64> <WORD 1-256> [ge <0-128>] [le <0-128>]
[id <1-2147483647>]
```

Use the same command to add additional prefixes to the list.

3. To rename the list:

```
ipv6 prefix-list <WORD 1-64> name <WORD 1-64>
```

4. Display the prefix list:

```
show ipv6 prefix-list <WORD 1-256> [vrf WORD < 1-16 >] [vrfids WORD < 1-512 >] [WORD < 1-64 >]
```

Example

Create an IPv6 prefix list:

```
Switch:1<config>#ipv6 prefix-list list4 2001:DB8::/32 ge 32 le 32
```

To rename the list:

Switch:1<config>#ipv6 prefix-list list4 name list5

Variable Definitions

Use the data in the following table to use the ipv6 prefix-list command.

Variable	Value
<word 1–64=""></word>	Specifies the IPv6 prefix-list name.
<word 1–256=""></word>	Specifies the IPv6 prefix and length.
ge <0–128>	Specifies the minimum length to match.
	Lower bound and higher bound mask lengths together can define a range of networks.

Variable	Value
id <1-2147483647>	Specifies the Prefix list ID.
le <0–128>	Specifies the maximum length to match.
	Lower bound and higher bound mask lengths together can define a range of networks.
name <word 1-64=""></word>	Names the prefix list. The default value is none.

Use the data in the following table to use the show ipv6 prefix-list command.

Variable	Value
<word 1–256=""></word>	Specifies the IPv6 prefix and length.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Configuring IP Route Policies

Configure a route policy so that the device can control routes that certain packets can take. For example, you can use a route policy to deny certain Border Gateway Protocol (BGP) routes.

The route policy defines the matching criteria and the actions taken if the policy matches.

About this task

After you create and enable the policy, you can apply it to an interface. You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In this case, all sequence numbers under the given policy apply to that filter.

Create and enable the policy for IS-IS accept policies for Fabric Connect for Layer 3 Virtual Services Networks (VSNs) and IP Shortcuts, then apply the IS-IS accept policy filters. For more information on IS-IS accept policy filters, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Note:

After you configure route-map in Global Configuration mode or VRF Router Configuration mode, the device enters Route-Map Configuration mode, where you configure the action the policy takes, and define other fields the policy enforces.

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.

Note:

You cannot configure IPv4 and IPv6 route-maps on the same match statement.

Procedure

1. Enter Route-Map Configuration mode:

```
enable
configure terminal
route-map WORD<1-64> <1-65535>
```

2. At the route-map prompt, define the match criteria for the policy:

```
match {as-path WORD < 0-256 > | community WORD < 0-256 > | community-exact enable | extcommunity WORD < 0-1027 > | interface WORD < 0-259 > | local-preference < 0-2147483647 > |metric < 0-65535 > | metric-type-isis < any | internal | external > | network WORD < 0-259 > | next-hop WORD < 0-259 > | protocol WORD < 0-60 > | route-source WORD < 0-259 > | route-type < any | | local | internal | external | external - 1 | external - 2 > | tag | WORD < 0-256 > | vrf | WORD < 1-16 > | vrfids | WORD < 0-512 > }
```

- 3. Define the action the policy takes:
 - a. Allow the route:

```
permit
```

OR

b. Ignore the route:

```
no permit
```

4. Define the set criteria for the policy:

```
set {as-path WORD < 0-256 > | as-path-mode <tag|preprend> | automatic-tag enable | community WORD < 0-256 > | community-mode <additive|none|unchanged> | injectlist WORD < 0-1027 > | ip-preference <0-255 > | local-preference <0-2147483647 > | mask <A.B.C.D> | metric <0-65535 > | metric-type <type1|type2> | metric-type-internal <0-1 > | metric-type-isis <none|internal|external> | metric-type-live-metric | next-hop WORD < 0-256 > | nssa-pbit enable | origin <igp | egp | incomplete > | origin-egp-as <0-65535 > | tag WORD < 0-256 > | weight <0-65535 > }
```

5. Display current information about the IP route policy:

```
show route-map [WORD<1-64>] [<1-65535>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Enter Route-Map Configuration mode. At the route-map prompt, define the fields the policy enforces. Define the action the policy takes. Display current information about the IP route policy.

Variable Definitions

Use the data in the following table to use the match command.

Variable	Value
as-path WORD<0-256>	Configures the device to match the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified AS-lists. This field is used only for BGP routes and ignored for all other route types.
	WORD <0-256> specifies the list IDs of up to four AS-lists, separated by a comma.
	Use the no operator to disable match as-path: no match as-path WORD<0-256>
community WORD<0-256>	Configures the device to match the community attribute of the BGP routes against the contents of the specified community lists. This field is used only for BGP routes and ignored for all other route types.
	WORD <0-256> specifies the list IDs of up to four defined community lists, separated by a comma.
	Use the no operator to disable match community: no match community WORD<0-256>
community-exact enable	When disabled, configures the device so match community- exact results in a match when the community attribute of the BGP routes match an entry of a community-list specified in match-community.
	When enabled, configures the device so match-community- exact results in a match when the community attribute of the BGP routes matches all of the entries of all the community lists specified in match-community.
	enable enables match community-exact.

Variable	Value
	Use the no operator to disable match community-exact: no match community-exact enable
extcommunity WORD <0-1027>	Configures the device to match the extended community.
	WORD<0-1027> specifies an integer value from 1–1027 that represents the community list ID you want to create or modify.
interface WORD <0-259>	If configured, configures the device to match the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other route types.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match-interface: no match interface WORD <0-259>
local-preference <0-2147483647>	Configures the device to match the local preference, applicable to all protocols.
	<0-2147483647> specifies the preference value.
metric <0-65535>	Configures the device to match the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored.
	<0-65535> specifies the metric value. The default is 0.
network WORD <0-259>	Configures the device to match the destination network against the contents of the specified prefix lists.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match network: no match network WORD <0-259>
next-hop WORD<0-259>	Configures the device to match the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match next hop: no match next-hop WORD<0-259>
protocol WORD<0-60>	Configures the device to match the protocol through which the route is learned.
	WORD <0-60> is xxx, where xxx is local, ospf, ebgp, ibgp,isis, rip, static, or a combination separated by ,
	Use the no operator to disable match protocol: no match protocol WORD<0-60>

Variable	Value
route-source WORD<0-259>	Configures the system to match the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match route source: no match route-source WORD<0-259>
route-type {any local internal external external-1 external-2}	Configures a specific route type to match (applies only to OSPF routes).
	any local internal external external-1 external-2 specifies OSPF routes of the specified type only (External-1 or External-2). Another value is ignored.
tag WORD<0-256>	Specifies a list of tags used during the match criteria process. Contains one or more tag values.
	WORD<0-256> is a value from 0-256.
[vrf WORD<1-16>] [vrfids WORD<0-512>]	Configures a specific VRF to match (applies only to RIP routes).

Use the data in the following table to use the set command.

Variable	Value
as-path WORD<0-256>	Configures the device to add the AS number of the AS-list to the BGP routes that match this policy.
	WORD<0-256> specifies the list ID of up to four defined AS-lists separated by a comma.
	Use the no operator to delete the AS number: no set aspath WORD<0-256>
as-path-mode <tag prepend></tag prepend>	Configures the AS path mode.
	Prepend is the default configuration. The device prepends the AS number of the AS-list specified in set-as-path to the old aspath attribute of the BGP routes that match this policy.
	ℜ Note:
	Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS.
automatic-tag enable	Configures the tag automatically. Used for BGP routes only.
	Use the no operator to disable the tag: no set automatic-tag enable

Variable	Value
community WORD<0-256>	Configures the device to add the community number of the community list to the BGP routes that match this policy.
	WORD <0-256> specifies the list ID of up to four defined community lists separated by a comma.
	Use the no operator to delete the community number: no set community WORD<0-256>
community-mode <additive none < td=""><td>Configures the community mode.</td></additive none <>	Configures the community mode.
unchanged>	additive—the device prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy.
	none—the device removes the community path attribute of the BGP routes that match this policy to the specified value.
injectlist WORD<0-1027>	Configures the device to replace the destination network of the route that matches this policy with the contents of the specified prefix list.
	WORD<0-1027> specifies one prefix list by name.
	Use the no operator to disable set injectlist: no set injectlist
ip-preference <0-255>	Configures the preference. This applies to accept policies only.
	<0-255> is the range you can assign to the routes.
local-preference <0-65535>	Configures the device to match the local preference, applicable to all protocols. <0–655356> specifies the preference value.
mask <a.b.c.d></a.b.c.d>	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.
	A.B.C.D is a valid contiguous IP mask.
	Use the no operator to disable set mask: no set mask
metric <0-65535>	Configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF for RIP, the original cost of the route or default-import-metric is used (applies to IS-IS routes also).
metric-type {type1 type2}	Configures the metric type for the routes to announce into the OSPF domain that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
metric-type-internal <0–1>	Configures the MED value for routes advertised to ebgp nbrs to the IGP metric value.
	<0-1> specifies the metric type internal.

Variable	Value
metric-type-isis <none external="" internal="" =""></none>	Configures the metric type for IS-IS routes. The default is none. This field is applicable only for IS_IS policies.
metric-type-live-metric	Configures the metric type for BGP routes. The default is disabled. This field is applicable only for BGP policies.
next-hop WORD <1-256>	Specifies the IP address of the next-hop router. Both IPv4 and IPv6 addresses are supported.
	Use the no operator to disable set next-hop: no set next-hop
nssa-pbit enable	Configures the not so stubby area (NSSA) translation P bit. Applicable to OSPF announce policies only.
	Use the no operator to disable set nssa-pbit: no set nssa-pbit enable
origin {igp egp incomplete}	Configures the device to change the origin path attribute of the BGP routes that match this policy to the specified value.
origin-egp-as <0-65535>	Indicates the remote autonomous system number. Applicable to BGP only.
tag <0-65535>	Configures the tag of the destination routing protocol. If not specified, the device forwards the tag value in the source routing protocol. A value of 0 indicates that this parameter is not configured.
	Note:
	This parameter is not supported on all hardware platforms.
weight <0-65535>	Configures the weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. Used for BGP only. A value of 0 indicates that this parameter is not configured.

Use the data in the following table to use the ${\tt name}$ command.

Variable	Value	
WORD<1-64>	Renames a policy and changes the name field for all	
	sequence numbers under the given policy.	

Job aid

Use the data in the following table to use the show route-map command output.

Table 13: Variable definitions

Variable	Value
NAME	Indicates the name of the route policy.
SEQ	Indicates the second index used to identify a specific policy within the route policy group (grouped by ID). Use this field to specify different match and set parameters and an action.
MODE	Indicates the action to take when this policy is selected for a specific route. Options are permit, deny, or continue. Permit indicates to allow the route. Deny indicates to ignore the route. Continue means continue checking the next match criteria configured in the next policy sequence; if none, take the default action in the given context.
EN	Indicates whether this policy is enabled. If disabled, the policy is not used.

IPv6 Basic Configuration using EDM

Use the procedures in this section to configure IPv6 basics using EDM.

Enabling the IPv6-mode boot config flag

Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits.

This flag is disabled by default. Use this procedure to enable (set to true) the IPv6-mode boot config flag.

When the IPv6-mode boot config flag is enabled, the maximum number of IPv4 routing table entries decreases. For scaling information, see Release Notes for VOSS.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit**,
- 2. Click Chassis.
- 3. Click the **Boot Config** tab.
- 4. Select the **EnableIpv6Mode** check box.
- 5. Click Apply.
- 6. Save the configuration, and then reboot the switch for the change to the IPv6-mode boot config flag to take effect.

Configuring IPv6 globally

Global configuration includes the following:

• IPv6 alternative routes: To avoid traffic interruption, enable alternative routes globally on the switch, to replace the best route with the next-best route if the best route becomes unavailable. By default, this feature is enabled.

Before you begin

Change the VRF instance as required to configure IPv6 globally on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the **Configuration > IPv6** folders.
- 2. Click IPv6.
- 3. Click the Globals tab.
- 4. In the **DefaultHopLimit** box, enter the preferred number of hops before packets drop.
- 5. To enable ICMP network unreachable messages, click **IcmpNetUnreach**.
- 6. In the **IcmpErrorInterval** box, enter the preferred interval for sending ICMPv6 error messages.
- 7. In the **IcmpErrorQuota** box, enter the preferred number of ICMP error messages that the system can send during the ICMP error interval.
- 8. To enable IPv6 multicast, click **IcmpMulticastRequestEnable**.
- 9. To enable IPv6 ICMP address unreachable messages, click IcmpAddrUnreach.
- To enable IPv6 ICMP port unreachable messages, click IcmpPortUnreach.
- 11. To enable IPv6 autoconfiguration, click Autoconfig.
- 12. To enable IPv6 static routes globally, click StaticRouteGlobalAdminEnabled.
- 13. To enable IPv6 Source Routing, click **SourceRouteEnable**.
- 14. To clear all IPv6 static routes, click **RouteStaticClear**.
- 15. To enable IPv6 alternative routes, click **AlternativeRouteEnable**.
- 16. To configure IPv6 ECMP globally, select enable or disable, in the EcmpEnable option box.
- 17. In the **EcmpMaxPath** box, enter the preferred number of equal-cost paths.
- 18. Click Apply.

Globals field descriptions

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

Name	Description
Forwarding	Configures whether this switch is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, the switch. Select forwarding for the switch to act as a router for IPv6 traffic. Select notForwarding to not act as a router for IPv6 traffic.
	The default is forwarding .
	You must enable forwarding to use Telnet or Ping with IPv6.
DefaultHopLimit	Configures the hop limit.
	The default is 64.
Interfaces	Shows the number of interfaces.
IfTableLastChange	Shows the date of the last interface table change.
IcmpNetUnreach	Enables ICMP network unreachable messages. The default is disabled.
IcmpErrorInterval	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000. An entry of 1 seconds results in no sent ICMPv6 error messages.
IcmpErrorQuota	Configures the number of ICMP error messages that the system can send during the ICMP error interval. A value of zero specifies not to send any.
	The default value is 50.
IcmpMulticastRequestEnable	Globally enables or disables the IPv6 ICMP echo multicast request feature. This is enabled by default.
IcmpAddrUnreach	Enables or disabled ICMP address unreachable messages.
	This is enabled by default.
IcmpPortUnreach	Enables or disables ICMP port unreachable messages.
	This is enabled by default.
Autoconfig	Enables or disables stateless address autoconfiguration.
	This is disabled by default.
StaticRouteGlobalAdminEnabled	Enables IPv6 static routes globally. If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM.
	The default is enabled.

Name	Description
SourceRouteEnable	Globally enables or disables the IPv6 Source Routing feature. It is disabled by default.
PrefixListTableSize	Displays the prefix list table size.
RoutePrefTableSize	Displays the route preference table size.
AlternativeRouteEnable	Globally enables or disables the IPv6 alternative route feature. By default, this feature is enabled.
EcmpEnable	Enables or disables the IPv6 ECMP globally. By default, it is disabled.
EcmpMaxPath	Globally configures the maximum number of ECMP paths.
	The number of paths supported is either 1 to 4 or 1 to 8, depending on your hardware platform.
	When ECMP is enabled, the default value is either 4 or 8 depending on your hardware platform.
	You cannot configure this feature unless ECMP is enabled globally.

Configuring an IPv6 interface

You must configure an IPv6 interface for a VLAN or brouter port before you can assign an IPv6 address to the interface.

Before you begin

 You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address. The switch supports port-based, protocol-based, and MAC-source-based VLANs.

For information about how to configure VLANs, see the following documents:

- Configuring VLANs, Spanning Tree, and NLB for VOSS
- Configuring Link Aggregation, MLT, SMLT and vIST for VOSS
- Change the VRF instance as required to configure an IPv6 interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

You can also configure an IPv6 interface for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can create both types of interfaces.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **IPv6** folders.
- 2. Click IPv6.
- 3. Click the Interfaces tab.

- 4. Click Insert.
- 5. In the Interface field, click Port or VLAN.
- 6. Select a port or VLAN.
- 7. Click OK.
- 8. Select the **AdminStatus** field to activate the interface.
- 9. Configure the remaining parameters as required.
- 10. Click Insert.
- 11. Click Apply.

Interfaces field descriptions

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

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	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 value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Туре	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500.
	Different hardware platforms support different MTU values.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).

Name	Description
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN.
	You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
	Different hardware platforms support different MAC offset ranges.
RSMLTEnable	Shows whether RSMLT is enabled on the interface.
	The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface.
	The default value is disabled (false).

Configuring an IPv6 brouter port interface

You must configure an IPv6 interface for a brouter port before you can assign an IPv6 address to the interface.

Procedure

- 1. In the Device Physical View, select a port.
- 2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
- 3. Click IPv6.
- 4. Click the IPv6 Interface tab.

- 5. Click Insert.
- 6. Enter the interface identifier.
- 7. Select the **AdminStatus** field to activate the interface.
- 8. Configure the remaining parameters as required.
- 9. Click Insert.
- 10. Click Apply.

Interfaces field descriptions

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

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	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 value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Туре	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500.
	Different hardware platforms support different MTU values.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.

Name	Description
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN. You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range. Different hardware platforms support different MAC offset ranges.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface.
	The default value is disabled (false).

Configuring an IPv6 VLAN interface

You must configure an IPv6 interface for a VLAN before you can assign an IPv6 address to the interface.

Before you begin

- You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address. The switch supports port-based, protocol-based, and MAC-source-based VLANs.
 - For information about how to configure VLANs, see the following documents:
 - Configuring VLANs, Spanning Tree, and NLB for VOSS
 - Configuring Link Aggregation, MLT, SMLT and vIST for VOSS

Procedure

- 1. In the navigation pane, expand the **Configuration** > **VLAN** folders.
- 2. Click VLANs.
- 3. Click the **Basic** tab.
- 4. Select a VLAN.
- 5. Click IPv6.
- 6. Click the IPv6 Interface tab.
- 7. Click Insert.
- 8. Enter the interface identifier.
- 9. Select the **AdminStatus** field to activate the interface.
- 10. Configure the remaining parameters as required.
- 11. Click Insert.
- 12. Click Apply.

IPv6 Interfaces field descriptions

Use the data in the following table to use the IPv6 Interfaces tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	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 value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Туре	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP

Name	Description
	addresses defined on this interface. The default value is 1500.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN.
	You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
	Different hardware platforms support different MAC offset ranges.
RSMLTEnable	Shows whether RSMLT is enabled on the interface.
	The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface.
	The default value is disabled (false).

Assigning IPv6 addresses to interfaces

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

You can assign an IPv6 address to a VLAN or brouter port.

Before you begin

Change the VRF instance as required to assign IPv6 addresses to interfaces on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6.
- 3. Click the **Addresses** tab.
- 4. Click Insert.
- 5. In the Interface field, click Port or VLAN.
- 6. Select the interface.
- 7. Click OK.
- 8. Type the IPv6 address and prefix length.
- 9. Click Insert.
- 10. Click Apply.

Addresses field descriptions

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

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	Specifies the IPv6 address to which this entry applies.
	1 Important:
	If the IPv6 address exceeds 116 octets, the object identifiers (OIDS) of instances of columns in this row are more than 128 subidentifiers and you cannot use SNMP-v1, SNMPv2c, or SNMPv3 to access them.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Туре	Specifies the type of address. The default is unicast.

Name	Description
Origin	Specifies the origin of the address. The following list shows the possible origins:
	• other
	manual
	• dhcp
	linklayer
	• random
Status	Specifies the status of the address, describing whether the address is used for communication. The following list shows the possible statuses:
	prefered (default)
	deprecated
	invalid
	inaccessible
	• unknown
	tentative
	duplicate
Created	Specifies the sysUpTime of the creation of this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
LastChanged	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
ValidLifetime	Shows how long, in seconds, the address is valid.
PrefLifetime	Shows how long, in seconds, the address is in use.

Assigning IPv6 addresses to a brouter port interface

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

You can assign an IPv6 address to a VLAN or brouter port.

Before you begin

Change the VRF instance as required to assign IPv6 addresses to a brouter port interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
- 2. Click IPv6.
- 3. Click the Ipv6 Addresses tab.
- 4. Click Insert.
- 5. Click Insert.
- 6. Click Apply.

IPv6 Addresses field descriptions

Use the data in the following table to use the IPv6 Addresses tab.

Name	Description
Interface	Specifies the port.
Addr	Specifies the IPv6 address to which this entry applies.
	Note:
	If the IPv6 address exceeds 116 octets, the object identifiers (OIDS) of instances of columns in this row are more than 128 subidentifiers and you cannot use SNMP-v1, SNMPv2c, or SNMPv3 to access them.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Туре	Specifies the type of address. The default is unicast.
Origin	Specifies the origin of the address.
Status	Specifies the status of the address, describing whether the address is used for communication.
Created	Specifies the sysUpTime of the creation of this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
LastChanged	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.

Assigning an IPv6 address to a VLAN

Assign an IPv6 address to a VLAN.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **VLAN** folders.
- 2. Click VLANs.
- 3. Click the **Basic** tab.
- 4. Select a VLAN interface.
- 5. Click IPv6.
- 6. Click the IPv6 Addresses tab.
- 7. Click Insert.
- 8. Type the IPv6 address and length in the fields.
- 9. Click Insert.
- 10. Click Apply.

IPv6 Addresses field descriptions

Use the data in the following table to use the IPv6 Addresses tab.

Name	Description
Interface	Identifies the address to which the address is assigned.
Addr	Specifies an IP address that is associated with a VLAN.
AddrLen	Specifies the prefix address length value for this address.
Туре	Specifies the type of address: either unicast or anycast.
Origin	Specifies the origin of the address as one of the following:
	• other
	• manual
	• dhcp
	linklayer
	• random
Status	Shows the status of the address and if it can be used for communication.

Name	Description
Created	Shows the time this address entry was created.
LastChanged	Shows the time this address entry was last updated.

Creating IPv6 static routes

Use static routes to manually configure routes to destination IPv6 address prefixes.

Before you begin

- Enable IPv6 forwarding.
- Change the VRF instance as required to create IPv6 static routes on a specific VRF instance.
 Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- Click IPv6.
- 3. Click the Globals tab.
- 4. Select the StaticRouteGlobalAdminEnabled check box.

If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM. The default is enabled.

- 5. Click Apply.
- 6. Click the Static Routes tab.
- 7. Click Insert.
- 8. In the **Dest** field, type the IPv6 address.
- 9. In the **PrefixLength** field, type the length of the prefix for the IPv6 address.
- 10. In the **NextHop** field, type the IPv6 address of the router through which the specified route is accessible.
- 11. Beside the Interface field, click Port or Vlan or Tunnel.
- 12. Select the interface, and then click **OK**.
- 13. In the **Cost** field, type a number for the distance.
- 14. Select the **Enable** check box.
- 15. Click Insert.

Static Routes field descriptions

Use the data in the following table to use the **Static Routes** tab.

Name	Description
Dest	Specifies the IPv6 destination network address. The prefix value must match the PrefixLength.
PrefixLength	Specifies the number bits you want to advertise from the prefix. The prefix value must match the value in the Dest field. The range is 0 to 128.
NextHop	Specifies the IPv6 address of the next hop on this route. You do not need to specify the next hop if the devices directly connect to one another. Configure the next hop if the two nodes do not share the same network prefix but reside on the same link.
Interface	Specifies the interface to which this entry applies. You must specify the port or VLAN if the next hop is a link-local address.
Cost	Specifies the cost or distance ratio to reach the destination for this node. The range is 1-65535. The default value is 1.
Enable	Enables the static route on the port. The default value is enable.
Status	Shows the status of the static route as one of the following:
	notReachable: The route is not reachable and no neighbor request entry is built to resolved the next- hop. This status appears if no route or neighbor exists to reach the next-hop of the static route.
	 tryToResolve: The route is not reachable but a neighbor request entry is built to resolve the next- hop. This status appears if a local equivalent route exists in the system to reach the next-hop but the neighbor is not learned.
	reachableNotInRtm: The static route is reachable but it is not in RTM. This status appears if the static route is reachable, but it is not the best among alternative static routes.
	 reachableInRtm: The static route is reachable and it is in RTM. This status appears if the static route is reachable, and it is the best among alternative static routes to be added into RTM.
Preference	Specifies the routing preference of the destination IPv6 address. The range is 1-255. The default value is 5.

Configuring IPv6 route preferences

Change IPv6 route preferences to force the routing protocols to prefer one route over another. Configure route preferences to override default route preferences and give preference to routes learned for a specific protocol.

Before you begin

Change the VRF instance as required to configure IPv6 route preferences on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task



Changing route preferences is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. If you want to change default preferences for routing protocols, do so before you enable the protocols.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6.
- 3. Click the RoutePref tab.
- 4. In the **ConfiguredValue** column, double-click a parameter to change the preference for the given protocol.
- 5. Click Apply.

RoutePref field descriptions

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

Name	Description
DefaultValue	Specifies the default preference value for the specified protocol.
Protocol	Specifies the protocol name.
ConfiguredValue	Configures the preference value for the specified protocol.

View Route Information

View routes information to view the current configuration.

About this task

IPv6 host routes created for the IPv6 local interfaces do not display in the routing table.

Before you begin

Change the VRF instance as required to view route information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Select IPv6.
- 3. Select the **Routes** tab.

Routes field descriptions

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

Name	Description
Dest	Specifies the IPv6 destination network address. The prefix value must match the PrefixLength.
PfxLength	Specifies the number bits you want to advertise from the prefix. The prefix value must match the value in the Dest field.
Index	Specifies the unique value that identifies the route among the routes to the same network layer destination.
Interface	Specifies the interface to which this entry applies.
NextHop	Specifies the IPv6 address of the next hop on this route.
Protocol	Specifies the routing protocol, such as OSPF.
Metric	Specifies the metric assigned to this interface. The default value of the metric is the reference bandwidth or ifSpeed. The value of the reference bandwidth is configured by the rcOspfv3ReferenceBandwidth object.
	For more information about reference bandwidth, see Globals field descriptions on page 192.
NextHopId	Identifier of the next-hop, hostname, or mac address.
Age	Specifies the number of seconds since the route was last updated or is last active.
Туре	Specifies the type of route.
PathType	Specifies the type of path.
SrcVrfld	Specifies the source VRF instance.
Pref	Specifies the preference.

Viewing IPv6 Default Routers

View the table of default routers learned from router advertisement messages.

A maximum of four routers are visible in the default routers list.



Not all hardware platforms include a dedicated, physical management interface. For more information about supported interfaces, see your hardware documentation.

Procedure

- 1. In the navigation tree, expand the **Configuration > IPv6** folders.
- 2. Click IPv6.
- Click the **Default Routers** tab.

Default Routers field descriptions

Use the data in the following table to use the Default Routers tab.

Name	Description
Address	Specifies the learned router address for an IPv6 default routers entry.
IfIndex	Specifies the interface number for an IPv6 default routers entry.
Lifetime	Specifies the remaining router lifetime.
Active	Specifies if the default router is active for an IPv6 default routers entry.

Configure a Circuitless IPv6 Interface

Before you begin

Change the VRF instance as required to configure a Circuitless IPv6 interface on a specific VRF instance.

About this task

You can use a circuitless IPv6 (CLIPv6) interface to provide uninterrupted connectivity to your system.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click IPv6.
- 3. Click the Circuitless IP tab.

- 4. Click Insert.
- 5. In the **Interface** field, assign a CLIP interface number.
- 6. Type the IPv6 address and prefix length.

Circuitless IPv6 field descriptions

Use the data in the following table to use the Circuitless IPv6 tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	Specifies the IPv6 address to which this entry applies.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.

Configuring IPv6 Prefix List

Use IPv6 prefix lists to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Before you begin

• Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

- 1. In the navigation pane, open the following folders: **Configuration > IPv6**.
- 2. Click Policy.
- 3. In the **Ipv6-Prefix List** tab, click **Insert**.
- 4. Edit the parameters as required.
- 5. Click Insert.

Ipv6-Prefix list field descriptions

Use the data in the following table to use the **lpv6–Prefix List** tab.

Name	Description
Id	Specifies the list identifier. The range is 1 to 2147483647.
Prefix	Specifies the prefix IPv6 address.
PrefixMaskLen	Specifies the length of the prefix mask. You must enter the full 128-bit mask to exact a full match of a specific IPv6 address (for example, when creating a policy to match the next-hop).

Name	Description
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name can be from 1 to 64 characters in length.
MaskLenFrom	Specifies the lower bound on the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Specifies the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configuring an IPv6 Route Policy

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
- 2. Click Policy.
- 3. Click the **Route Policy** tab.
- 4. Click Insert.
- 5. Enter the appropriate information for your configuration in the Insert Route Policy dialog box.
- 6. Click Insert.

Route Policy Field Descriptions

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

Name	Description	
Id	Specifies the ID of an entry in the Prefix list table.	
SequenceNumber	Specifies a policy within a route policy group.	
Name	Specifies the name of the policy. This command changes the name field for all sequence numbers under the given policy.	
Enable	Indicates whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored. The default is disabled.	
Mode	Specifies the action to take when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route). The default is permit.	
MatchProtocol	Selects the appropriate protocol. If configured, matches the protocol through which the route is learned. This field is used	

Name	Description	
	only for RIP Announce purposes. The default is to enable all match protocols.	
MatchNetwork	Specifies if the system matches the destination network against the contents of the specified prefix list.	
MatchipRouteSource	Specifies if the system matches the next-hop IP address for RIF routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.	
	Click the ellipsis button and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key.	
	You can also change this field in the Route Policy tab of the Policy dialog box.	
MatchipRouteDest	Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.	
MatchNextHop	Specifies if the system matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.	
	Click the ellipsis button and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.	
MatchInterface	Specifies if the system matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route.	
	Click the ellipsis button and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.	
MatchRouteType	Configures a specific route type to match (applies only to OSPF routes).	
	Externaltype1 and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.	
MatchMetric	Specifies if the system matches the metric of the incoming advertisement or existing route against the specified value (1 to 65535). If 0, this field is ignored. The default is 0.	
MatchMetricTypeIsis	Specifies the match metric type field in the incoming ISIS routes in accept policy.	

Name	Description	
MatchAsPath	Configures if the system matches the BGP autonomus system path. Applicable to BGP only. This overrides the BGP neighbor filter list information.	
MatchCommunity	Filters incoming and outgoing updates based on a Community List. Applicable to BGP only. The default is disable.	
MatchCommunityExact	Indicates if the match must be exact (that is, all of the communities specified in the path must match). Applicable to BGP only. The default is disabled.	
MatchTag	Specifies a list of tags used during the match criteria process. Applicable to BGP only. It contains one or more tag values.	
MatchVrf	Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).	
MatchLocalPref	Specifies the local preference value to be matched.	
NssaPbit	Configures or resets the P bit in specified type 7 link state advertisements (LSA). By default, the Pbit is always configured in case the user configures the Pbit to a disable state for a particular route policy other than all type 7. LSAs associated with that route policy have the Pbit cleared. With this intact, not so stubby area (NSSA) area border router (ABR) does not perform translation of these LSAs to type 5. The default is enable.	
SetRoutePreference	Configures a value from 0 to 255. The default value is 0. If the default is configured, the global preference value is used.	
	When configured to a value greater than zero, specifies the route preference value assigned to the routes that matches the policy. This feature applies to accept policies only.	
SetMetricTypeInternal	Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The value must be 0 or 1. The default is 0.	
SetMetricTypelsis	Sets the metric type IS-IS.	
SetMetric	Configures the system to use the metric value for the route while announcing or redistributing. The default-import-metric is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or the default value is used (applies to IS-IS routes also). The default is 0.	
SetMetricType	Configures the metric type for the routes to announce into the OSPF routing protocol that matches this policy. Applicable to OSPF protocol only. The default is type 2. This field is applicable only for OSPF announce policies. The default is type2.	
SetNextHop	Configures the IP address of the next-hop router. Applicable to BGP only. The default is 0.0.0.0.	

Name	Description	
SetInjectNetList	Configures the destination network of the route that matches this policy with the contents of the specified prefix list. Click the ellipsis button and choose from the list in the Set Inject NetList dialog box.	
SetMask	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.	
SetAsPath	Indicates the AS path value to use whether the SetAsPathMode field is Tag or Prepend. Applicable to BGP only.	
	Note:	
	Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS .	
SetAsPathMode	Configures if the system converts the tag of a route into an AS path. Applicable to BGP protocol only. The mode is either Tag or Prepend tag. The value is applicable only while redistributing routes to BGP The default is prepend.	
	Note:	
	Prepend is not applicable to an iBGP peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS .	
SetAutomaticTag	Enables the automatic tag feature. Applicable to BGP protocol only. The default is disable.	
SetCommunityNumber	Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.	
SetCommunityMode	Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged.	
	Unchanged—keeps the community attribute in the route path as it is.	
	None—removes the community in the route path additive.	
	Append—adds the community number specified in SetCommunityNumber to the community list attribute.	
SetExtCommunity	Configures a BGP community. The values are 0 to 256.	
SetExtCommunityMode	Configures the extended-community mode. The value can be append, unchanged, or overwrite. The default value is unchanged.	
	append — creates another community string.	
	unchanged — keeps the community attribute as it is.	

Name	Description	
	overwrite — changes the current value.	
SetOrigin	Configures the origin for the BGP protocol to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.	
SetLocalPref	Configures the local preference for the BGP protocol only. The system uses this value during the route decision process for the BGP protocol. The default is 0.	
SetOriginEgpAs	Indicates the remote autonomous system number for the BGP protocol. The default is 0.	
SetWeight	Configures the weight value for the routing table for the BGP protocol. This field must be used with the match as-path condition. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.	
SetTag	Configures the list of tags used during the match criteria process for the BGP protocol. The default is 0.	
Ipv6SetNextHop	Specifies the address of the IPv6 next hop router.	

Configuring IPv6 Route Redistribution Policies

About this task

Configure route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF or BGP. Use a route policy to control the redistribution of routes. You can redistribute routes for the Global router (VRF 0) and within a user-defined VRF but not between different VRFs.

Before you begin

- · Ensure 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 > IPv6**.
- 2. Click Policy.
- 3. Click the Route Redistribution tab.
- 4. Click Insert.
- 5. Choose the protocol and route source.
- 6. Select **Enable**.
- 7. Choose the route policy to apply to the redistributed routes.

- 8. Configure other parameters as required.
- 9. Click Insert.

Route Redistribution Field Descriptions

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

Name	Description
DstVrfld	Specifies the destination VRF ID.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrfld	Specifies the source VRF ID.
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. The default is disabled.
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. The default is 0.
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. The default is type2.

Chapter 4: Neighbor discovery

This chapter provides concepts and procedures to complete IPv6 neighbor discovery configuration.

Neighbor discovery

IPv6 nodes (routers and hosts) on the same link use neighbor discovery (ND) to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services for IPv4 with the Address Resolution Protocol (ARP) and router discovery. In IPv6 ND performs a function similar to ARP (Address Resolution Protocol) in IPv4.

Hosts use ND to discover the routers in the network that you can use as the default routers, and to determine the link-layer address of neighbors attached to local links. Routers also use ND to discover neighbors and link-layer information. ND updates the neighbor database with valid entries, invalid entries, and entries migrated to various locations.

The ND protocol provides the following services:

· address and prefix discovery

Hosts determine the set of addresses that are on-link for the given link. Nodes determine which addresses or prefixes are locally reachable or remote with address and prefix discovery.

router discovery

Hosts discover neighboring routers with router discovery. Hosts establish neighbors as default packet-forwarding routers.

parameter discovery

Hosts and routers discover link parameters such as the link MTU or the hop-limit value placed in outgoing packets.

· address autoconfiguration

Hosts configure an address for an interface with address autoconfiguration.

· duplicate address detection

Hosts and nodes determine if an address is assigned to another router or a host.

· address resolution

Hosts determine link-layer addresses (MAC for Ethernet) of the local neighbors (attached on the local network), provided the IP address is known.

next-hop determination

Hosts determine how to forward local or remote traffic with next-hop determination. The next hop can be a local or remote router.

neighbor unreachability detection

Hosts determine if the neighbor is unreachable, and if address resolution must be performed again to update the database. For neighbors you use as routers, hosts attempt to forward traffic through alternative default routers.

redirect

Routers inform the host of more efficient routes with redirect messages.

Neighbor discovery uses three components:

host-router discovery

Host-router discovery performs the following functions:

- router discovery
- prefix discovery
- parameter discovery
- address autoconfiguration
- · host-host communication

Host-host communication performs the following functions:

- address resolution
- next-hop determination
- neighbor unreachability detection
- duplicate address detection
- · route redirect

Note:

When a neighbor transitions to the STALE state, to initiate Neighbor Unreachability detection (NUD), a duplicate copy of the traffic destined to this neighbor is sent to the switch Control Processor (CP) on a low priority queue (queue 0). The original packet is forwarded to this neighbor. Once NUD is initiated, the hardware records are updated and the traffic is no longer sent to the CP. When a high rate of such traffic is sent to CP, the switch can drop some of these packets due to the in built CP rate limiting feature, which protects the CP from DOS attacks.

Use the command show qos cosq-stats cpu-port to view drop statistics on the CPU queue. This design does not result in loss of traffic.

Use the command ipv6 nd reachable-time <0-3600000> to increase the default value of 3000 milliseconds which in turn delays the scenario of data path sending STALE neighbor destined packets to the CP.

Recommended values for reachable time is 180000, and for retransmit interval is 5000.

ND messages

The following table compares the ICMP message types.

Table 14: IPv4 and IPv6 neighbor comparison

IPv4 Function	IPv6 Function	Description
ARP request message	Neighbor solicitation message	A node sends this message to determine the link-layer address of a neighbor or to verify that a neighbor is still reachable through a cached link-layer address. You can also use neighbor solicitations for duplicate address detection.
ARP reply message	Neighbor advertisement	A node sends this message either in response to a received neighbor solicitation message or to communicate a link-layer address change.
ARP cache	Neighbor cache	The neighbor cache contains information about neighbor types on the network.
Gratuitous ARP	Duplicate address detection New VRRP master interface announcement	A host or node sends a request with its own IP address to determine if another router or host uses the address. If the sender receives a reply, then there is a device with a duplicate address. Both hosts and routers use this function. Gratuitous ARP can also be used to announce the new VRRP master interface so that all
Router solicitation message (optional)	Router solicitation message (required)	switches can adjust their MAC tables. The host sends this message after it detects a change in a network interface operational state. The message includes a request for routers to generate router advertisement immediately rather than at the scheduled time.
Router advertisement message (optional)	Router advertisement message (required)	Routers send this message to advertise their presence with various links and Internet parameters either periodically or in response to a router solicitation message. Router advertisements contain prefixes that you use for on-link determination or address configuration, and a suggested hop limit value.
Redirect message	Redirect message	Routers send this message to inform hosts of a better first hop for a destination.

Neighbor discovery cache

The neighbor discovery cache lists information about neighbors in the network and can contain the following types of neighbors:

• static: a configured neighbor

· local: a device on the local system

· dynamic: a discovered neighbor

The following table describes the states in the neighbor cache.

Table 15: Neighbor cache states

State	Description
Incomplete	Address resolution is in progress and the system has not yet determined the link-layer address of the neighbor.
	The neighbor cache may also enter the Incomplete state when the switch cannot confirm subsequent reachability during the ND process for router neighbors. By contrast, the system deletes host neighbors, rather than enter the Incomplete state, if ND fails to confirm reachability.
	① Tip:
	Router neighbors: when the R bit is set in the received neighbor advertisement
	Host neighbors: when the R bit is not set in the received neighbor advertisement
Reachable	A node receives positive confirmation within the last reachable time period.
Stale	Reachability of the neighbor is unknown.
	Until the system sends traffic to the neighbor, make no attempt to verify its reachability.
Delay	A time period longer than the reachable time period passes since the node received the last positive confirmation, and a packet was sent within the last DELAY_FIRST_PROBE_TIME period.
	If no reachability confirmation is received within the DELAY_FIRST_PROBE_TIME period after entering the DELAY state, neighbor solicitation is sent and the state changes to probe.
Probe	Reachability confirmation is sought from the device every retransmit timer period.

The following events involve Layer 2 and Layer 3 interaction during processing and affect the neighbor cache:

- · flushing the virtual LAN (VLAN) MAC
- · removing a VLAN or brouter port
- · performing an action on all VLANs
- removing a port from a VLAN
- removing a port from a spanning tree group (STG)
- removing a multilink trunk (MLT) group from a VLAN
- removing an MLT port from a VLAN
- removing an MLT port from an STG
- performing an action that disables a VLAN, such as removing all ports from a VLAN
- disabling a tagged port that is a member of multiple routable VLANs

Router discovery

IPv6 nodes discover routers on the local link with router discovery.

IPv6 nodes discover routers on the local link with router discovery.

Router advertisement

Configured interfaces on an IPv6 router send router-advertisement messages. Interfaces also send router advertisements in response to router-solicitation messages from IPv6 nodes on the link.

Router solicitation

An IPv6 host without a configured unicast address sends router solicitation messages.

Host autoconfiguration

The switch can automatically configure a host (node), and assign IPv6 addresses automatically. This process is called stateless address autoconfiguration (SLAAC). The neighbor discovery (ND) protocol performs autoconfiguration.

Stateless autoconfiguration enables serverless basic configuration of IPv6 nodes and renumbering from a mathematical perspective.

Stateless autoconfiguration uses the following equation:

Stateless autoconfiguration = network prefix (router advertisement) + IPv6 interface identifiers

Stateless autoconfiguration uses the network prefix information in the router advertisement and integrates this with the interface ID to form the node global address(es).

■ Note:

The switch cannot autoconfigure an IPv6 address local to itself because IPv6 routers do not process router advertisements in the same manner as hosts. That is, routers check only for consistency in information advertised in IPv6 Router Advertisements on the same link.

Tip:

You must manually assign all addresses/prefixes local to the switch.

Assuming an EUI-64 based interface ID is used, the IPv6 interface address is created from the 48bit (6-byte) MAC address as follows:

- 1. EUI-64 hexadecimal digits 0xff-fe are inserted between the third and fourth bytes of the MAC address to obtain the EUI-64.
- 2. The universal or local bit, the second lower-order bit of the first byte of the MAC address, is complemented.

For example, the IPv6 identifier for host A uses the MAC address 00-AA-00-3F-2A-1C.

To automatically assign an address, the following occurs:

Convert to EUI-64 format

00-AA-00-FF-FE-3F-2A-1C

2. Complement the Universal/Local (U/L) bit.

The first byte in binary form is 00000000. When the seventh bit is complemented, it becomes 00000010 (0x02).

The result is 02-AA-00-FF-FE-3F-2A-1C or 2AA:FF:FE3F:2A1C.

Host A with MAC address 00-AA-00-3F-2A-1C, combined with network prefix 2001::/64 provided by router advertisement, uses an IPv6 address 2001::2AA:FF:FE3F:2A1C.

A host generates a link-local address with the prefix FE80 regardless of whether an IPv6 router is present or not.

The link-local address for a node with the MAC address 00-AA-00-3F-2A-1C is FE80::2AA:FF:FE3F: 2A1C.

The following list explains the states of an autoconfiguration address:

- Tentative: the address is being verified as unique (link-local address)
- Valid: an address from which unicast traffic can be sent and received; can be in one of two states—either preferred or deprecated
- Preferred: an address for which uniqueness was verified for unrestricted use; can be in one of three states—either tentative, preferred, or deprecated
- Deprecated: an address that remains valid but is withheld for new communication

Invalid: an address for which a node can no longer send or receive unicast traffic

Neighbor Discovery Configuration using CLI

Use the procedures in this section for neighbor discovery configuration using CLI.

Configuring an IPv6 discovery prefix

Configure the discovery prefixes to send in router advertisement.

About this task

Hosts on the link use router advertisements to perform IPv6 autoconfiguration.

The discovery prefix controls which IPv6 addresses will be automatically configured, and for how long they are valid.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Create neighbor discovery prefixes for an interface:

```
ipv6 nd prefix-interface WORD<0-255> [eui <1-3>] [no-advertise] [no-autoconfig] [no-onlink]
```

3. Modify an existing neighbor discovery prefix:

```
ipv6 nd prefix WORD<0-255> infinite [no-advertise] [preferred-life <0-4294967295>] [valid-life <0-4294967295>]
```

Example

Create a new neighbor discovery prefix:

Switch:1(config-if)#ipv6 nd prefix-interface fd48:bfb6:4c09:9499::1/64

Variable definitions

Use the data in the following table to use the ipv6 nd prefix and ipv6 nd prefix-interface commands.

Variable	Value
eui <1–3>	Configures the EUI address. The values are:
	• (1) EUI not used
	(2) EUI with Universal/Local bit (U/L) complement enabled
	• (3) EUI used without U/L
	Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global— and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the prefix length is 64 or less. The default is EUI not used.
	If you select EUI not used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. IF you select either EUI used with UL complement or EUI used without UL complement, an associated IPv6 adress is created by concatenating the specified prefix with the EUI-64 interface ID.
infinite	Configures the prefix valid lifetime so it never expires. The default is disabled, which means the prefix expires.
no-advertise	Removes the prefix from the neighbor advertisement. The default is disabled, which means the prefix is advertised.
no-autoconfig	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. The value is a 1-bit flag. The default is enabled.
no-onlink	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. The value is a 1-bit flag. The default is enabled.
preferred-life <0-4294967295>	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.
	The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.

Variable	Value
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
valid-life <0-0-4294967295>	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.
	A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
WORD <0-255>	Specifies the IPv6 address and prefix.

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring route advertisement

Configure route advertisement in IPv6 for neighbor discovery (ND).

About this task

IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

Procedure

1. Enter Interface Configuration mode:

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



If the 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 number of neighbor solicitation messages from duplicate address detection:

```
ipv6 nd dad-ns <0-600>
```

3. Configure the hop limit sent in router advertisements:

```
ipv6 nd hop-limit <0-255>
```

4. Enable managed address configuration (M-bit) on the router:

```
ipv6 nd managed-config-flag
```

5. Configure the MTU for router advertisements:

```
ipv6 nd mtu <0-9500>
```

6. Enable other stateful configuration (O-bit) on the router:

```
ipv6 nd other-config-flag
```

7. Configure the router lifetime included in router advertisement:

```
ipv6 nd ra-lifetime <0-9000>
```

8. Configure the neighbor reachable time:

```
ipv6 nd reachable-time <0-3600000>
```

9. Configure the time between neighbor solicitation messages:

```
ipv6 nd retransmit-timer <0-4294967295>
```

10. Configure the maximum time allowed between sending unsolicited multicast router advertisements:

```
ipv6 nd rtr-advert-max-interval <4-1800>
```

11. Configure the minimum time allowed between sending unsolicited multicast router advertisements:

```
ipv6 nd rtr-advert-min-interval <3-1350>
```

12. Enable periodic router advertisement messages:

```
ipv6 nd send-ra
```

Example

Configure the maximum time between sending unsolicited router advertisements:

Switch:1(config-if)#ipv6 nd rtr-advert-max-interval 700

Configure the minimum time between sending unsolicited router advertisements:

Switch:1(config-if) #ipv6 nd rtr-advert-min-interval 500

Enable periodic router advertisement messages:

Switch:1(config-if)#ipv6 nd send-ra

Variable definitions

Use the data in the following table to use the ipv6 nd commands.

Variable	Value
dad-ns <0-600>	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD).
	A value of 0 disables the DAD process on this interface.
	A value of 1 sends one advertisement without retransmissions.
hop-limit <0–255>	Specifies the current hop limit field sent in router advertisements from this interface.
	The value must be the current diameter of the Internet.
	A value of zero indicates that the advertisement does not specify a hop-limit value.
	The default is 64.
managed-config-flag	Enables the system to configure the M-bit, or managed address configuration flag, in the router advertisements
	When set, the M-bit flag indicates that addresses are available through DHCPv6.
	If the M flag is set, the O flag is redundant because DHCPv6 returns all available configuration information.
	If neither the M nor O flags are set, no information is available through DHCPv6.
	The default is disabled.
mtu <0-9500>	Shows the MTU value sent in router advertisements on this interface.

A value of zero indicates that the system sends no MTU options. The default is 0. Other-config-flag Enables the O-bit, or other stateful configuration, flag in the router advertisement. Other stateful configuration autoconfigures received information without addresses. When set, the O-bit flag indicates that other configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network. If neither the M nor O flags are set, no information is available through DHCPv6. The default is disabled. 7 Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. 7 Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Tetransmit-timer <0-4294967295> Tetransmit-timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. The default is 0.	Variable	Value
other-config-flag Enables the O-bit, or other stateful configuration, flag in the router advertisement. Other stateful configuration autoconfigures received information without addresses. When set, the O-bit flag indicates that other configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network. If neither the M nor O flags are set, no information is available through DHCPv6. The default is disabled. ra-lifetime <0-9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0-3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> retransmit-timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		•
in the router advertisement. Other stateful configuration autoconfigures received information without addresses. When set, the O-bit flag indicates that other configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network. If neither the M nor O flags are set, no information is available through DHCPv6. The default is disabled. ra-lifetime <0-9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0-3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> retransmit-timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system).		The default is 0.
information without addresses. When set, the O-bit flag indicates that other configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network. If neither the M nor O flags are set, no information is available through DHCPv6. The default is disabled. ra-lifetime <0–9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0–3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system).	other-config-flag	Enables the O-bit, or other stateful configuration, flag in the router advertisement.
configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network. If neither the M nor O flags are set, no information is available through DHCPv6. The default is disabled. ra-lifetime <0–9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0–3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		
available through DHCPv6. The default is disabled. ra-lifetime <0-9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0-3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		configuration information is available through DHCPv6; for example, DNS-related information or
ra-lifetime <0–9000> Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. reachable-time <0–3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		
router advertisements sent from this interface. This value must be either 0, or 4 to 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800. Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		The default is disabled.
A value of zero indicates that the system is not a default router. The default is 1800. Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.	ra-lifetime <0–9000>	
default router. The default is 1800. reachable-time <0–3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. retransmit-timer <0-4294967295> Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		This value must be either 0, or 4 to 9000 seconds.
reachable-time <0–3600000> Specifies a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		•
router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		The default is 1800.
Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.	reachable-time <0-3600000>	
node is considered reachable after a reachability confirmation event. The default is 0. Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		The value zero means unspecified (by this system).
retransmit-timer <0-4294967295> Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		node is considered reachable after a reachability
retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur.		The default is 0.
The value configures the amount of time the system waits for the transmission to occur.	retransmit-timer <0-4294967295>	retransmit timer field in the router advertisement
waits for the transmission to occur.		The value zero means unspecified (by this system).
The default is 0.		
		The default is 0.
rtr-advert-max-interval <4–1800> Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface.	rtr-advert-max-interval <4–1800>	
The default is 600.		The default is 600.

Variable	Value
rtr-advert-min-interval <3–1350>	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface.
	The default is 200.
send-ra	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface.
	The default is enabled.

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

System interface values versus advertised values

There are differences in the relationship between the system interface values and advertised values related to Neighbor Disovery (ND). The information in this section describes differences and similarities and provides examples for important IPv6 interface and IPv6 ND commands.

Comparison of default values per interface and as advertised

The following table compares the default behavior of values per interface and advertised values.

Default values per interface	Default advertised values
hop-limit 64	hop-limit 64
mtu 1500	mtu 0 (unspecified)
reachable-time 30000 ms	reachable-time 0 (unspecified)
retransmit-timer 1000 ms	retransmit-timer 0 (unspecified)

What happens when you change the per interface value and the advertised value?

When you change per-interface values from default to non-default values, the system changes the advertised values to match the interface values.

For example, when you enter the ipv6 interface mtu 1300 command the values become

- interface mtu 1300
- advertised mtu 1300

Then, when you enter the show ipv6 nd interface command, the system marks the mtu value with an (i) which signifies that the ND advertised value is inherited from the interface configuration.

Example: Changing both values

Switch:1(config-if) #ipv6 interface mtu 1300

```
Switch:1(config-if)#
show ipv6 nd interface GigabitEthernet 1/1
```

```
Port Ipv6 Nd

IFID BTR RTR- MAX- MIN- LIFE- MANAG OTHER DAD-NS MTU HOP REACH- RETRANS- ADV INT INT TIME LIMIT TIME TIME ABLE MIT FLAG CONF

320 1/1 True 600 200 0 False False 1 1300(i) 64(d) 0(d) 0(d) Note: (s) = Set, (d) = Default, (i) = inherit
```

What happens when you change the per interface value but do not change the advertised value?

To change the per-interface value from the default value to a non-default value but retain the advertised value of 0 (unspecified), you must enter two commands.

For example, to set the reachable-time to 60000 but retain the advertised value of the reachable-time parameter at 0, enter the following commands:

```
ipv6 interface reachable-time 6000
ipv6 nd reachable-time 0
```

When you enter the show ipv6 nd interface command, the system marks the reachable-time value with an (s) to signify that this value is explicitly set by the ND configuration.

Example: Changing only the per interface value

```
Switch:1(config-if) #ipv6 interface reachable-time 60000
Switch:1(config-if) #ipv6 nd reachable-time 0
```

```
Switch:1(config-if)#
show ipv6 nd interface GigabitEthernet 1/1
```

```
Port Ipv6 Nd

IFID BTR RTR- MAX- MIN- LIFETIME MANAG OTHER DAD-NS MTU HOP REACH- RETRANS- ADV INT INT CONF LIMIT TIME TIME ABLE MIT FLAG

320 1/1 True 600 200 0 False False 1 1500(i) 64(d) 0(s) 0(d) Note: (s) = Set, (d) = Default, (i) = inherit
```

Configuring the neighbor cache

Configure the address translation table used to map IPv6 addresses to physical addresses. You can manually add static neighbors to the cache.



IPv6 static neighbors are not supported on SMLT.

About this task

Neighbor cache in IPv6 is similar to the IPv4 Address Resolution Protocol (ARP) table.

The neighbor cache is a set of entries for individual neighbors to which traffic was recently sent.

You make entries on the neighbor on-link unicast IP address, including information such as the linklayer address.

A neighbor cache entry contains information used by the Neighbor Unreachability Detection algorithm, including the reachability state, the number of unanswered probes, and the time the next Neighbor Unreachability Detection event is scheduled.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a static neighbor:

```
ipv6 neighbor WORD<0-128> port {slot/port[sub-port]} mac
<0x00:0x00:0x00:0x00:0x00:0x00> [vlan <1-4059> ]
```

When you create a static neighbor, it always remains in the reachable state. This differs from the general neighbor cache behavior where, among other things, timers and neighbor unreachability detection events can be generated.

Example

Create a static neighbor:

```
Switch:1(config) #ipv6 neighbor 3000::3 port 1/11 mac 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the ipv6 neighbor command.

Variable	Value
mac <0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the MAC address.
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized,

Variable	Value
	you must also specify the sub-port in the format slot/port/sub-port.
vlan <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.
WORD<0-128>	Specifies the IPv6 address in hexadecimal colon format.

Viewing cached destination information

View the destination cache to see next-hop addresses for destinations.

The destination cache is only populated or updated when IPv6 packets originate locally on the central processor of the switch.

The main purpose of the destination cache is to store, on a per-destination basis, the dynamic Path MTU value currently used when transmitting packets from the local system to the remote destination.

The system uses the PMTU value to calculate how many bytes can fit into an individual packet before fragmentation should be applied.

About this task

The command output shows the following information:

- the IPv6 destination address
- the IPv6 address for the next hop to the destination
- the path maximum transmission unit (MTU) for the destination
- the time, in seconds, since an ICMPv6 packet-too-big message was received

Not all parameters are available in non-default VRFs.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. View the destination cache for all interfaces:

```
show ipv6 dcache [vrf WORD<1-16> | vrfids WORD<0-512>]
```

3. View the destination cache for a brouter port:

```
show ipv6 dcache gigabitethernet { slot/port[/sub-port] }
```

4. View the destination cache for a management port:

show ipv6 dcache mgmtethernet mgmt



Note:

This step only applies to hardware with a dedicated, physical management interface.

5. View the destination cache for a specific tunnel ID:

```
show ipv6 dcache tunnel <1-2000>
```

6. View the destination cache for a VLAN:

```
show ipv6 dcache vlan <1-4059>
```

7. Clear the destination cache:

```
clear ipv6 dcache [qiqabitethernet {slot/port[/sub-port]}]
[mgmtethernet {slot/port[/sub-port]}][tunnel <1-2000>][vlan
<1-4059>] [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example

Switch:1(config-if)#show ipv6 do	cache			
	IPv6 Destination Cache Informat			
Destination Address PMTU PMTU_AGE	NEXT HOP	VID/BID/TID		
ff02:0:0:0:0:0:1 1500 0	0:0:0:0:0:0:0	V-22	real -	
1 out of 1 Total Num of Destinaton Cache Entries displayed.				

Variable definitions

Use the data in the following table to use the show ipv6 dcache and clear ipv6 dcache commands.

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.
{slot/port[/sub-port]}	Identifies a single slot and port. If the 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 name.
vrfids WORD<0-512>	Specifies the VRF ID.

Neighbor Discovery Configuration using EDM

Use the procedures in this section for neighbor discovery configuration using EDM.

Configuring an IPv6 discovery prefix

Configure the discovery prefixes to send in router advertisement.

Before you begin

Change the VRF instance as required to configure an IPv6 discovery prefix on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

You can also configure an IPv6 interface for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6.
- 3. Click the **Discovery Prefix** tab.
- 4. Click Insert.
- 5. Beside the Interface field, click Port or VLAN.
- 6. Select a port or VLAN.
- 7. Click OK.
- 8. Specify the prefix and prefix length.
- 9. Click Insert.
- 10. Click Apply.

Discovery Prefix field descriptions

Use the data in the following table to use the Discovery Prefix tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.

Name	Description
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.
VLanId	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.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
ValidLifetime	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.
	A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
PreferedLifetime	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.
	The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.

Name	Description
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both gloal and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used.
	If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. IF you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 adress is created by concatenating the specified prefix with the EUI-64 interface ID.
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configuring an IPv6 discovery prefix port

Configure the discovery prefixes to send in router advertisement.

About this task

Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
- 2. Click IPv6.
- 3. Click Discovery Prefix.
- 4. Click Insert.
- 5. Click OK.
- 6. Specify the prefix and prefix length.
- 7. Click Insert.
- 8. Click Apply.

IPv6 Discovery Prefix field descriptions

Use the data in the following table to use the IPv6 Discovery Prefix tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.
VLanId	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.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
ValidLifetime	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.
	A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
PreferedLifetime	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.
	The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current

Name	Description
	preferred lifetime, you must lower the preferred lifetime value first.
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both gloal and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used.
	If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. IF you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 adress is created by concatenating the specified prefix with the EUI-64 interface ID.
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configuring an IPv6 discovery prefix on a VLAN

Configure the discovery prefixes to send in router advertisement.

About this task

Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

Procedure

- 1. In the navigation pane, expand the following folders: **VLAN** > **VLANs**.
- 2. Click the Basic tab.
- 3. Select an interface row.
- 4. Click IPv6.

- 5. Click the **IPv6 Discovery Prefix** tab.
- 6. Click Insert.
- 7. Specify the prefix and prefix length.
- 8. Configure the remaining parameters, as required.
- 9. Click Insert.
- 10. Click Apply.

IPv6 Discovery Prefix field descriptions

Use the data in the following table to use the IPv6 Discovery Prefix tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.
VLanid	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.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
ValidLifetime	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.
	A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.

Name	Description
PreferedLifetime	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.
	The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.
	The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both gloal and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used.
	If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. IF you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 adress is created by concatenating the specified prefix with the EUI-64 interface ID.
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configuring route advertisement

Configure route advertisement in IPv6 for neighbor discovery (ND).

Before you begin

Change the VRF instance as required to configure route advertisement on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

Note:

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

You can also configure an IPv6 interface for a brouter port through the Edit > Port > IPv6 navigation path, and for a VLAN through the VLAN > VLANs > Basic > IPv6 navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6.
- Click the Route Advertisement tab.
- 4. Double-click a parameter to change the current value. You cannot modify the parameters in gray shading.
- 5. Click Apply.

Route Advertisement field descriptions

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
MaxInterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on

Name	Description
	this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Configuring route advertisement on an IPv6 interface for a brouter port

Configure route advertisement in IPv6 for neighbor discovery (ND).

About this task

IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

Note:

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

You can also configure an IPv6 interface for a VLAN through the VLAN > VLANs > Basic > IPv6 navigation path.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
- 2. Click IPv6.
- Click the Route Advertisement tab.
- 4. Double-click a parameter to change the current value.

You cannot modify the parameters in gray shading.

5. Click Apply.

Route Advertisement field descriptions

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
Maxinterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.

Name	Description
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Configuring route advertisement on an IPv6 interface for a VLAN

Configure route advertisement in IPv6 for neighbor discovery (ND).

About this task

IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.



™ Note:

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

Procedure

- 1. In the navigation pane, expand the following folders: **VLAN** > **VLANs**.
- 2. Click the **Basic** tab.
- 3. Select an interface row, and click **IPv6**.
- 4. Click Route Advertisement.
- 5. Double-click a parameter to change the current value. You cannot modify the parameters in gray shading.
- 6. Click Apply.

Route Advertisement field descriptions

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
Maxinterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.

Name	Description
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Configuring the neighbor cache

Configure the address translation table used to map IPv6 addresses to physical addresses. You can manually add static neighbors to the cache.



Note:

IPv6 static neighbors are not supported on SMLT.

Before you begin

Change the VRF instance as required to configure neighbor cache on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

Neighbor cache in IPv6 is similar to the IPv4 Address Resolution Protocol (ARP) table. The neighbor cache is a set of entries for individual neighbors to which traffic was sent recently. You make entries on the neighbor on-link unicast IP address, including information such as the link-layer address. A neighbor cache entry contains information used by the Neighbor Unreachability Detection algorithm, including the reachability state, the number of unanswered probes, and the time the next Neighbor Unreachability Detection event is scheduled.

Procedure

- 1. In the navigation tree, expand the following folders: Configuration > IPv6.
- 2. Click IPv6.
- 3. Click the **Neighbors** tab.
- 4. Click Insert.
- 5. Beside the Interface field, click Port or Port in Vlan.
- 6. Select a port or VLAN.
- 7. Configure the remaining parameters as required.
- 8. Click Insert.
- 9. Click **Apply**.

Neighbors field descriptions

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

Name	Description
Interface	Specifies the interface to which this entry applies.
NetAddress	Specifies the IP address of the media-dependent physical address.
PhyAddress	Specifies the MAC address.
Interface	Specifies a physical port ID or a MLT port ID.

Name	Description
LastUpdated	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
Туре	Specifies the mapping type from manually configured entries. While the selection of either dynamic, static, or local is allowed; static is currently the only valid selection.
State	Specifies the Neighbor Unreachability Detection state for the interface after the address mapping in this entry is used. If Neighbor Unreachability Detection is not in use (for example, for IPv4), this object is always unknown. The options include the following:
	reachable: confirmed reachability
	stale: unconfirmed reachability
	delay: waiting for reachability confirmation before entering the probe state
	probe: actively probing
	invalid: an invalidated mapping
	unknown: state cannot be determined.
	incomplete: address resolution is being performed
ВМас	Specifies the backbone MAC address.
Cvid	Specifies the customer VID.

Viewing cached destination information

View the destination cache to see next-hop addresses for destinations.

The destination cache is only populated or updated when IPv6 packets are locally originated on the central processor of the switch.

The main purpose of the destination cache is to store, on a per-destination basis, the dynamic Path MTU value currently used when transmitting packets from the local system to the remote destination. The PMTU value itself is used to calculate how many bytes can fit into an individual packet before fragmentation should be applied.

Before you begin

Change the VRF instance as required to view cached destination information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click IPv6.
- 3. Click the **Destination Cache** tab.

Destination Cache field descriptions

Use the data in the following table to use the **Destination Cache** tab.

Name	Description
DestAddr	Shows the IPv6 destination address.
Interface	Shows the interface number that is used to reach the destination.
NextHop	Shows the IPv6 address for the next hop to the destination.
IfType	Specifies the interface type (tunnel, VLAN, or brouter) or virtual circuit (VRRP, RSMLT).
IfData	Displays additional information about virtual circuits. For instance, for a VRRP or RSMLT the virtual router ID displays. If the interface type is tunnel, VLAN, or brouter, no additional information displays.
Pmtu	Shows the path maximum transmission unit (MTU) for the destination.
PmtuAge	Shows the time, in seconds, since an ICMPv6 packet too big message was received.

Chapter 5: DHCP Relay

Table 16: Dynamic Host Configuration Protocol Relay for IPv6 product support

Feature	Product	Release introduced	
For configuration details, see Config	For configuration details, see Configuring IPv6 Routing for VOSS.		
IPv6 DHCP Relay	VSP 4450 Series	VOSS 4.1	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 4.1	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 6.2	
	XA1400 Series	Not Supported	

DHCP Relay

The Dynamic Host Configuration Protocol (DHCP) for IPv6 (RFC 3315) enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. DHCP supports automatic allocation of reusable network addresses and of additional configuration parameters. This protocol is a stateful counterpart to stateless address autoconfiguration, and you can use it separately or concurrently with the latter to obtain configuration parameters. For more information about stateless address autoconfiguration, see Host autoconfiguration on page 106.

To request the assignment of one or more IPv6 addresses, a client first locates a DHCP server, and then requests the assignment of addresses and other configuration information from the server:

- 1. The client sends a solicit message to the All_DHCP_Relay_Agents_and_Servers (FF02::1:2) multicast address to find available DHCP servers.
- 2. Any server that can meet the requirements responds with an advertise message.
- 3. The client then chooses one of the servers and sends a request message to the server asking for confirmed assignment of addresses and other configuration information.
- 4. The server responds with a reply message that contains the confirmed addresses and configuration.

If a DHCP client does not need a DHCP server to assign it an IPv6 address, the client can obtain configuration information such as a list of available DNS servers or NTP servers through a single message and reply exchanged with a DHCP server.

IPv6 DHCP clients use link-local addresses to send and receive DHCP messages. To permit a DHCP client to send a message to a DHCP server that is not attached to the same link, you must configure a DHCP relay agent on the client link to relay messages between the client and server. The operation of the relay agent is transparent to the client.

A relay agent relays messages from clients and messages from other relay agents. The switch supports DHCP Relay for IPv6. Configure at least one relay agent when the client and server are in different networks.

You must configure the relay agent to use a list of destination addresses for available DHCP servers. The software does not support IPv6 multicast for site-local and global addresses.

The DHCP relay can be a Virtual Router Redundancy Protocol (VRRP) Address. The relay forwards the DHCP messages only if VRRP is in the Master state, otherwise the relay discards the messages.



Note:

Since DHCP cannot work on the backup VRRP if the master fails, to achieve optimum results and to leverage redundancy you must configure DHCP on the backup VRRP.

Clients listen for DHCP messages on UDP port 546. Servers and relay agents listen for DHCP messages on UDP port 547.

Remote ID

IPv6 DHCP Relay supports the remote ID parameter (RFC4649). After you enable remote ID on the switch, the relay agent adds information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The server can use the supplied information in the process of assigning the addresses, delegated prefixes, and configuration parameters that the client is to receive.

The remote ID option contains two fields:

- vendor ID
- · MAC address of the client

The switch uses a vendor ID of 1584.

Limitations

The following list identifies configuration limitations:

- You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.
- The maximum number of servers to which a relay can send a message from one client, is 10.
- You can configure the number of forwarding paths per system. For information on the maximum limit, see Release Notes for VOSS.

DHCP Relay Network Topology and Workflow

The following example depicts the interaction between a DHCP client and a server:

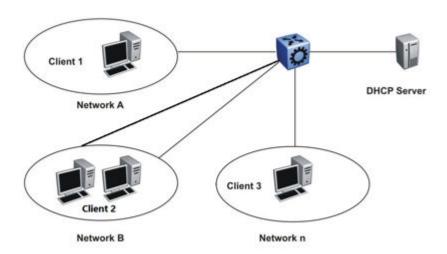


Figure 9: DHCP Client-Relay-Server Architecture

The following list outlines the operations that the DHCP relay agent performs to forward the message to the server :

• When a client sends a request for the IP address or configuration parameters, the server responds with the details as requested by the client.

Note:

There should be at least one relay agent when client and server are located in different networks.

• A DHCP Relay IPv6 is established only between agents within the context of each VRF and when no cross VRF interaction is present.

Note:

All_DHCP_Servers multicast address option is not implemented for IPv6, as there is no IPv6 MCAST support for site-local and global address.

DHCP Relay Configuration using CLI

Use the procedures in this section to configure DHCP Relay using CLI.

Configuring a DHCP Relay forwarding path

Configure a forwarding path to specify the relay agent address and the DHCP server address to which to forward packets.

To use DHCP Relay for IPv6, you must configure at least one forwarding path and enable the relay on one interface.

Before you begin

For a VRF other than GlobalRouter, the interface must be first associated to that VRF.

About this task

The relay agent can use the IPv6 address of the interface or the VRRP global address linked to that interface. The relay forwards the DHCP messages only if VRRP is in the master state, otherwise the relay discards the messages.

You can configure only one relay agent on an interface. If you need to change the relay agent, you must delete all the forwarding paths with the old relay agent, and then configure the new relay agent.

For scaling information on DHCP Relay forwarding paths, see Release Notes for VOSS.

Procedure

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

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

2. Configure a forwarding path:

```
ipv6 dhcp-relay fwd-path WORD<0-255> WORD<0-255> [enable]
```

If you configure the forwarding path globally, the relay agent address can be any configured IP address of the relay interface or the VRRP global address linked to the relay interface.

3. To configure a forwarding path on an interface, enter Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

OR

```
interface vlan <1-4059>
```

4. Configure a forwarding path:

```
ipv6 dhcp-relay fwd-path WORD<0-255> [enable] [vrid WORD<1-255>]
```

If you configure the forwarding path on an interface, the relay agent address is either the smallest IP configured on the interface or the first VRRP global address configured, if the

relay is the VRRP master. You do not specify the relay agent address as part of the command.



Note:

IPv6 DHCP Relay is established only between agents within the context of each VRF.

Example

Configure a forwarding path globally:

Switch:1(config) #ipv6 dhcp-relay fwd-path 1111::1111 1234::1234 enable

Configure a forwarding path on an interface:

Switch:1(config) #interface GigabitEthernet 1/1 Switch:1(config-if) #ipv6 dhcp-relay fwd-path 1234::1234 enable

OR

Configure the VRRP master as the relay:

Switch:1(config-if) #ipv6 dhcp-relay fwd-path 1234::1234 vrid 12 enable

Variable definitions

Use the data in the following table to use the ipv6 dhcp-relay fwd-path 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.
enable	Enables the forwarding path. The default is disabled.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vrid WORD<1-255>	Specifies the VRRP ID to use the VRRP master as the relay agent interface.
WORD<0-255>	Specifies the IPv6 address of the DHCP server for the interface configuration.
WORD<0-255> WORD<0-255>	Specifies the IPv6 address of the relay agent interface and the IPv6 address of the DHCP server for the global configuration.

Configuring DHCP Relay for an interface

Configure the DHCP relay behavior on the interface.

About this task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Enable DHCP on the interface:

```
ipv6 dhcp-relay
```

3. Configure the maximum hop count:

```
ipv6 dhcp-relay max-hop <1-32>
```

4. Enable the remote ID:

```
ipv6 dhcp-relay remote-id
```

Example

Configure the maximum hop count:

```
Switch:1(config-if) #ipv6 dhcp-relay max-hop 30
```

Disable the remote ID:

```
Switch:1(config-if) #no ipv6 dhcp-relay remote-id
```

Variable definitions

Use the data in the following table to use the ipv6 dhcp-relay command.

Variable	Value
max-hop <1–32>	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
remote-id	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing DHCP Relay information

View DHCP Relay information to display the current configuration for the forwarding path and the interface configuration.

About this task

Not all parameters are available in non-default VRFs.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View DHCP Relay global information:

```
show ipv6 dhcp-relay {counters [vrf WORD<1-16> | vrfids WORD<0-512>] | fwd-path [vrf WORD<1-16> | vrfids WORD<0-512>]
```

3. View IPv6 DHCP Relay interface configuration:

show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[/subport] [-slot/port[/sub-port]] [,...]} | vlan <1-4059>}



Note:

The no ipv6 dhcp-relay command disables DHCP on the interface but does not delete the entry.

Example

Variable definitions

Use the information in the following table to help you use the show ipv6 dhcp-relay command.

Variable	Value
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan<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.
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

DHCP Relay Configuration using EDM

Use the procedures in this section to configure DHCP Relay using EDM.

Configuring a DHCP Relay forwarding path

Configure a forwarding path to specify the relay agent address and the DHCP server address to which to forward packets.

To use DHCP Relay for IPv6, you must configure at least one forwarding path and enable the relay on one interface.

Before you begin

Change the VRF instance as required to configure a DHCP Relay forwarding path on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

The relay agent can use the IPv6 address of the interface or the VRRP global address linked to that interface. The relay forwards the DHCP messages only if VRRP is in the Master state, otherwise the relay discards the messages.

You can configure only one relay agent on an interface. If you need to change the relay agent, you must delete all the forwarding paths with the old relay agent, and then configure the new relay agent.

For scaling information on DHCP Relay forwarding paths, see Release Notes for VOSS.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click **DHCP Relay**.
- 3. Click the Forward Path tab.
- 4. Click Insert.
- 5. In the **AgentAddr** field, type the address of the input interface that forwards the packets.
- 6. In the **ServerAddr** field, type the address of the DHCP server.
- 7. Select Enabled.
- 8. Click Insert.

Forward Path field descriptions

Use the data in the following table to use the **Forward Path** tab.

Name	Description
AgentAddr	Specifies the IP address of the input interface (relay agent) on which the DHCP request packets are received for forwarding. This address is the IPv6 or VRRP global address of either a brouter port or a VLAN for which forwarding is enabled.
ServerAddr	Specifies the IP address of the DHCP server. The request is unicast to the server address.
Enabled	Enables DHCP Relay for the system. The default is disabled (clear).

Configuring DHCP Relay for an interface

Configure the DHCP relay behavior on the interface.

Before you begin

Change the VRF instance as required to configure DHCP Relay for an interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

You can modify the DHCP Relay configuration for a brouter port through the **Edit** > **Port** > **IPv6** navigation path, and for a VLAN through the **VLAN** > **VLANs** > **Basic** > **IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click DHCP Relay.
- 3. Click the **Interface** tab.
- 4. Click Insert.
- 5. Beside the **IfIndex** field, click **Port** or **Vlan**.
- 6. Select a port or VLAN, and then click **OK**.
- 7. Click Insert.

Interface field descriptions

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

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface. For the brouter port, the value is the ifindex of the port and, in the case of the VLAN, the value is the ifindex of the VLAN.
МахНор	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteldEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).

Modifying DHCP Relay for a VLAN

Modify the existing DHCP relay behavior on the VLAN interface.

About this task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs.
- Click the Basic tab.
- 4. Select a VLAN.
- 5. Click IPv6.
- 6. Click the **DHCP Relay** tab.
- 7. Double-click a cell to change the value.
- 8. Click Apply.

DHCP field descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface.
МахНор	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteldEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).
DhcpEnabled	Enables (true) or disables (false) DHCP Relay for an interface with an existing DHCP Relay configuration.

Modifying DHCP Relay for a port

Modify the existing DHCP relay behavior on the brouter port interface.

About this task

The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

- 1. In the Device Physical View, select the port.
- 2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.

- 3. Click IPv6.
- 4. Click **DHCP Relay**.
- 5. Double-click a cell to change the value.
- 6. Click Apply.

DHCP Relay field descriptions

Use the data in the following table to use the DHCP Relay tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface. For the brouter port, the value is the ifindex of the port and, in the case of the VLAN, the value is the ifindex of the VLAN.
МахНор	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteldEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).
DhcpEnabled	Enables (true) or disables (false) DHCP Relay for an interface with an existing DHCP Relay configuration. This field appears on the DHCP Relay tab for a brouter port only if you modify an existing configuration. This field does not appear if you create a new DHCP Relay port configuration.

Chapter 6: Tunneling

Table 17: IPv4 in IPv6 tunnels product support

Feature	Product	Release introduced			
For configuration details, see Config	For configuration details, see Configuring IPv6 Routing for VOSS.				
IPv4 in IPv6 tunnels	VSP 4450 Series	VOSS 4.1			
	VSP 4900 Series	VOSS 8.1			
	VSP 7200 Series	VOSS 4.2.1			
	VSP 7400 Series	VOSS 8.0			
	VSP 8200 Series	VOSS 4.1			
	VSP 8400 Series	VOSS 4.2			
_	VSP 8600 Series	VSP 8600 6.2			
	XA1400 Series	Not Supported			

Tunneling

Tunneling provides a mechanism to transfer IPv6 traffic through an IPv4—only network.

How tunneling works

IPv6 tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure.

At the tunnel source, or head end, the system encapsulates an IPv6 packet into an IPv4 packet and sends it to the remote tunnel destination.

The tunnel destination strips the IPv4 packet header and forwards the original IPv6 packet further into an IPv6 cloud.

These types of tunnels are called dual-stack tunnels because they support both IPv4 and IPv6.

Manually configured tunnels

Manually configured tunnels can provide communication between two isolated IPv6 domains over an IPv4 network.

Manually configured tunnels are point-to-point.

You can configure tunnel endpoints to create a point-to-point connection between two isolated IPv6 domains by configuring IPv6 and IPv4 addresses at each end of the tunnel.



Note:

The router or host at the source and destination ends of the tunnel must support both IPv4 and IPv6 protocol stacks.



Caution:

Ensure that all single-homed point-to-point traffic ingresses and egresses a configured tunnel. Otherwise the traffic is dropped.

IPv6 reachability enables tunnel forwarding but tunnel operational status depends on the IPv4 reachability of the tunnel endpoint.

The IPv4 tunnel endpoint configuration must be symmetrical; that is, if you configure a tunnel with a source of 10.10.10.1 and a destination of 11.11.11.1 from switch A, then Switch B must have a source of 11.11.11.1 and a destination of 10.10.10.1.

Tunnel interfaces are logical point-to-point interfaces.

You can enable dynamic routing when you enable a routing protocol, for example OSPFv3, on the tunnel interfaces.

Unicast routing protocols can detect link loss and redirect IPv6 route information

There is no explicit signaling protocol applied to IPv6-in-IPv4 configured tunnels (refer to RFC 4213).

Therefore, if the remote endpoint of a tunnel that terminates several Layer 3 hops away in the network fails, the local state of the tunnel remains active even though the endpoint has failed.

However, you can enable unicast routing protocols over tunnels, for example OSPFv3. These unicast routing protocols introduce their own protocol-specific signaling and, when a unicast routing protocol is present over the tunnel link, the routing protocol can detect link loss and re-direct the IPv6 route information to use an alternate, reachable nexthop.

Operational events that trigger tunnel state transition

The switch must be able to locally detect operational events that can trigger a tunnel state transition.

These events include:

- · deletion of local IPv4 interface
- change or loss of the IPv4 route to the remote tunnel endpoint
- change in the nexthop of the IPv4 route to the remote tunnel endpoint
- loss of the ARP entry for the nexthop router that is used to reach the IPv4 tunnel endpoint

Tunnels and MTU

You cannot configure the MTU for tunnels.

The default MTU value for tunnels is 1280.

Packets are forwarded through the tunnel using the line card network processing units (NPUs) only. Since the packets are not forwarded through the central processing unit (CPU) they do not impact the CPU load.

Tunnels and BGP+

You must configure an IPv6 tunnel and static routes on BGP+ peers when you use BGP+. For more information on IPv6 tunnel configuration for BGP+, see <u>Configuring BGP Services for VOSS</u>.

Limitations

The following list identifies tunnel configuration limitations.

- You cannot configure IPv6 CLIP addresses for IPv6-in-IPv4 tunnels. Also, you cannot configure an IPv6 CLIP interface as the source or destination endpoint of an tunnel.
- You cannot configure SMLT on the switch terminating a tunnel.
- Termination of tunnels on vIST peers is not supported.

Tunneling Configuration using CLI

Use the procedures in this section to configure Tunnel using CLI.

Configuring a tunnel

Configure a tunnel for IPv6 VLANs or brouter ports to communicate through an IPv4–only network. Create a point-to-point connection between the two isolated IPv6 devices by configuring the tunnel endpoints.

Do not create tunnels in a native IPv6 network.

Before you begin

 The router or host at the source and destination of the tunnel must support both IPv4 and IPv6 protocol stacks.

About this task

Manual tunnels are point-to-point, so you configure both source and destination addresses. You must configure both IPv6 and IPv4 addresses for both source and destination devices. The IPv6 addresses must represent the same network, for example 6666::1/96 and 6666::2/96.

Tunnel interfaces are automatically configured with a link-local address in the format fe80::<local_ipv4_source_address>.

You cannot configure the maximum transmission unit (MTU) for tunnels. The default MTU value for tunnels is 1280.

Procedure

1. Enter Global Configuration mode:

enable

configure terminal

2. Create a tunnel:

```
ipv6 tunnel <1-2000> source {A.B.C.D} address WORD<0-46> destination {A.B.C.D}
```

Example

Create tunnel 2:

```
Switch:1(config) #ipv6 tunnel 2 source 11.11.11.1 address 3000:0:0:0:0:0:0:1/64 destination 12.12.12.2
```

Variable definitions

Use the data in the following table to use the ipv6 tunnel command.

Variable	Value
<1-2000>	Configures the ID for the tunnel.
address WORD<0-46>	Assigns an IPv6 address and prefix to the tunnel.
destination {A.B.C.D}	Configures the address of the remote endpoint of the tunnel.
source {A.B.C.D}	Configures the address of the local endpoint of the tunnel.

Viewing tunnel interfaces

View tunnel interfaces to verify the current configuration and operational status of IPv6 tunnels.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Show IPv6 tunnel information:

```
show ipv6 tunnel [<1-2000>] [detail] [local {A.B.C.D}] [remote {A.B.C.D}]
```

Example

```
Switch:1#show ipv6 tunnel detail

Tunnel Interface Information

LOCAL ADDRESS REMOTE ADDRESS OPER STATUS TYPE

2 211.1.55.2 44.1.55.1 active manual
1 211.1.55.2 44.1.55.43 active manual
210 211.1.60.2 47.1.60.1 active manual
3 out of 3 Total number of entries displayed.
```

	======		=======		
	Address	Informa 	ation 		
IPV6 ADDRESS/PREFIX LENGTH			TYPE	ORIGIN	STATUS
43:210:0:0:0:0:0:2/64 fe80:0:0:0:0:0:0:d301:3702/64 44:211:0:0:0:0:0:2/64			UNICAST UNICAST	LINKLAYER MANUAL	PREFERRED
fe80:0:0:0:0:0:0:0301:3702/64				LINKLAYER	

Variable definitions

Use the data in the following table to use the **show ipv6 tunnel** command.

Variable	Value
<1–2000>	Shows information for a specific tunnel ID.
detail	Shows detailed address information for the tunnel.
local {A.B.C.D}	Shows information for a specific local address (the local endpoint of the tunnel).
remote {A.B.C.D}	Shows information for a specific remote address (the remote endpoint of the tunnel).

Modifying tunnel hop limits

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

About this task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Modify the hop limit:

```
ipv6 tunnel <1-2000> hop-limit <0-255>
```

Example

Modify the hop limit for tunnel ID 5:

```
Switch:1(config) #ipv6 tunnel 5 hop-limit 200
```

Variable definitions

Use the data in the following table to use the ipv6 tunnel command.

Variable	Value
<0–255>	Configures the maximum number of hops in the tunnel. The default value is 255.
<1–2000>	Specifies the tunnel ID.

Tunneling Configuration using EDM

Use the procedures in this section to configure Tunnel using EDM.

Configuring a tunnel

Configure a tunnel for IPv6 VLANs or brouter ports to communicate through an IPv4—only network. Create a point-to-point connection between the two isolated IPv6 devices by configuring the tunnel endpoints.

Do not create tunnels in a native IPv6 network.

Before you begin

 The router or host at the source and destination of the tunnel must support both IPv4 and IPv6 protocol stacks.

About this task

Manual tunnels are point-to-point, so you configure both source and destination addresses. You must configure both IPv6 and IPv4 addresses for both source and destination devices. The IPv6 addresses must represent the same network, for example 6666::1/96 and 6666::2/96.

Tunnel interfaces are automatically configured with a link-local address in the format fe80::<local_ipv4_source_address>.

You cannot configure the maximum transmission unit (MTU) for tunnels. The default MTU value for tunnels is 1280.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click Tunnel.
- 3. Click the **Tunnel Config** tab.
- 4. Click Insert.
- 5. Beside the **LocalAddress** field, click the button, and then select the IPv4 address for the local VLAN or brouter port.
- 6. In the **RemoteAddress** field, type the IPv4 address for the destination VLAN or brouter port.
- 7. In the **ID** field, type a number to represent the tunnel.

- 8. In the IPv6AddressAddr field, type the IPv6 address for the tunnel VLAN or brouter port.
- 9. In the **IPv6AddressPrefixLength** field, type the number of bits to advertise in the IPv6 address.
- 10. Click Insert.

Tunnel Config field descriptions

Use the data in the following table to use the **Tunnel Config** tab.

Name	Description
AddressType	Shows the address type over which the tunnel encapsulates packets.
LocalAddress	Configures the address of the local endpoint of the tunnel.
RemoteAddress	Configures the address of the remote endpoint of the tunnel.
EncapsMethod	Configures the tunnel mode, which is manual for manually configured tunnels.
ID	Configures the ID for the tunnel.
IfIndex	Shows the value of ifIndex that corresponds to the tunnel interface. A value of 0 indicates that the interface index has not yet been assigned. This field appears only on the Tunnel Config tab.
lpv6AddressAddr	Specifies the IPv6 address for the local VLAN or brouter port. This field appears only on the Insert Tunnel Config dialog box.
Ipv6AddressPrefixLength	Specifies the number of bits to advertise in the IPv6 address. This field appears only on the Insert Tunnel Config dialog box.

Modifying tunnel hop limits

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

Use this procedure to modify the hop limits for multiple tunnels simultaneously.

About this task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click Tunnel.
- 3. Click the Tunnel Interface tab.

- 4. Double-click the **HopLimit** value to modify the information as required.
- 5. Click **Apply**.

Tunnel Interface field descriptions

Use the data in the following table to use the **Tunnel Interface** tab.

Name	Description
Index	Identifies the tunnel interface internally. The value is derived from the tunnel ID.
EncapsMethod	Displays the encapsulation method for the tunnel: manual for manually configured tunnels and 6to4 for automatically configured tunnels.
HopLimit	Configures the maximum number of hops in the tunnel. The default value is 255.
Security	Indicates the type of security on the tunnel interface.
TOS	Displays the method used to configure the high 6 bits (the differentiated services codepoint) of the IPv4 type of service (TOS) or IPv6 traffic class in the outer IP header.
	A value of -1 indicates that the bits are copied from the payload header. A value of -2 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB module. A value from 0 to 63 indicates that the bit field is configured to the indicated value.
FlowLabel	Displays the method used to configure the IPv6 flow label value. This object is not required where AddressType indicates the tunnel is not over IPv6. A value of -1 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB. Any other value indicates that the flow label field is configured to the indicated value.
AddressType	Displays manual for a manually configured tunnel, or sixToFour for autoconfigured tunnels.
LocalInetAddress	Identifies the local endpoint address of the tunnel.
RemoteInetAddress	Identifies the remote endpoint of the tunnel.
EncapsLimit	Displays the maximum number of additional encapsulations permitted for packets undergoing encapsulation at this node. A value of -1 indicates that no limit exists, except as a result of the packet size.

Modifying tunnel hop limits for a specific tunnel

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

Use this procedure to modify the hop limits for a specific tunnel interface.

About this task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click Tunnel.
- 3. Click the **Tunnel Config** tab.
- 4. Select the tunnel row.
- 5. Click Tunnel Interface.
- 6. Double-click the **HopLimit** value to modify the information as required.
- 7. Click Apply.

Tunnel Interface field descriptions

Use the data in the following table to use the **Tunnel Interface** tab.

Name	Description
Index	Identifies the tunnel interface internally. The value is derived from the tunnel ID.
EncapsMethod	Displays the encapsulation method for the tunnel: manual for manually configured tunnels and 6to4 for automatically configured tunnels.
HopLimit	Configures the maximum number of hops in the tunnel. The default value is 255.
Security	Indicates the type of security on the tunnel interface.
TOS	Displays the method used to configure the high 6 bits (the differentiated services codepoint) of the IPv4 type of service (TOS) or IPv6 traffic class in the outer IP header.
	A value of -1 indicates that the bits are copied from the payload header. A value of -2 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB module. A value from 0 to 63 indicates that the bit field is configured to the indicated value.

Name	Description
FlowLabel	Displays the method used to configure the IPv6 flow label value. This object is not required where AddressType indicates the tunnel is not over IPv6. A value of -1 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB. Any other value indicates that the flow label field is configured to the indicated value.
AddressType	Displays manual for a manually configured tunnel, or sixToFour for autoconfigured tunnels.
LocalInetAddress	Identifies the local endpoint address of the tunnel.
RemotelnetAddress	Identifies the remote endpoint of the tunnel.
EncapsLimit	Displays the maximum number of additional encapsulations permitted for packets undergoing encapsulation at this node. A value of -1 indicates that no limit exists, except as a result of the packet size.

Viewing IPv6 addresses on a tunnel

View a tunnel for IPv6 addresses.

You can assign an IPv6 address to a VLAN or brouter port.

About this task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

You can also assign an IPv6 address through the **Edit** > **Port** > **IPv6** navigation path, and through the **VLAN** > **VLANs** > **Basic** > **IPv6** navigation path.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
- 2. Click Tunnel.
- 3. Click the **Tunnel Config** tab.
- 4. Click IPv6 Address.

Tunnel Config field descriptions

Use the data in the following table to use the **Tunnel Config** tab.

Name	Description
AddressType	Shows the address type over which the tunnel encapsulates packets.
LocalAddress	Configures the address of the local endpoint of the tunnel.
RemoteAddress	Configures the address of the remote endpoint of the tunnel.
EncapsMethod	Configures the tunnel mode, which is manual for manually configured tunnels.
ID	Configures the ID for the tunnel.
IfIndex	Shows the value of ifIndex that corresponds to the tunnel interface. A value of 0 indicates that the interface index has not yet been assigned. This field appears only on the Tunnel Config tab.

Chapter 7: OSPFv3

Table 18: OSPFv3 product support

Feature	Product	Release introduced
For configuration details, see Config	guring OSPF and RIP for VOSS.	
OSPFv3	VSP 4450 Series	VOSS 4.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 4.1
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported

OSPFv3

The Open Shortest Path First Protocol (OSPF) for IPv6, defined in RFC 2740 and RFC 5340, is an Interior Gateway Protocol used to distribute IPv6 routing information within a single Autonomous System (AS).

The IPv4 terms subnet and network are replaced in IPv6 by link. An IPv6 link is a communication medium between nodes at the link layer. You can assign multiple IP subnets (prefixes) to a link. Two IPv6 nodes with common or different prefixes can communicate over a single link.

OSPF for IPv6 operates on each link rather than each subnet as in IPv4. IPv6 makes the following changes to how packets are received and to the contents of network LSAs and hello packets:

- The OSPF packet contains no IPv6 addresses. LSA payloads carried in link state update packets contain IPv6 addresses.
- The following IDs remain at 32-bits and are not assigned IPv6 addresses: area IDs, LSA link state IDs, and OSPF router IDs.
- IPv6 OSPF neighbors use Router IDs to identify neighboring routers on broadcast and nonbroadcast multiaccess (NBMA) networks and for other communication media, point to point.

Flooding scope

LSA flooding scope is generalized in OSPFv3 and coded in the LS type field of the LSA. The following three flooding scopes are available for LSAs:

- Link scope: The LSA is not flooded beyond the local link.
- Area scope: The LSA is flooded in a single OSPF area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs.
- AS scope: The LSA is flooded through the routing domain. AS scope is used for ASexternal-LSAs.

Link-local addresses

IPv6 uses link-local addresses on a single link. Link-local addresses facilitate features such as neighbor discovery and autoconfiguration. Datagrams with link-local sources are not forwarded. Instead, routers assign link-local unicast addresses from the IPv6 address range.

OSPF for IPv6 does not assign link-local unicast addresses to physical segments attached to a router, it assumes that each router already has link-local unicast addresses assigned. The source for all OSPF packets sent on OSPF physical interfaces is the associated link-local unicast address. Routers learn link-local addresses for all other nodes on links. The nexthop information during packet forwarding includes the learned addresses.

OSPFv3 packets always use link-local addresses as the source and destination, except on a virtual link. All OSPFv3 packets sent over a virtual link use global addresses.

Link LSA is the only OSPF LSA type that includes link-local addresses. Link-local addresses must not be advertised in other LSA types.

Authentication

OSPFv3 for IPv6 requires the IP authentication header and the IP encapsulating security payload for authentication and security. OSPFv3 does not support the authentication feature from OSPFv2.

Packet format

OSPFv3 runs directly over IPv6. All other addressing information is absent in OSPF packet headers. OSPFv3 is network-protocol-independent. LSA types contain addressing information.

OSPFv3 implements the following packet changes from OSPFv2:

- The hello packet and database description packet operations fields are expanded to 24 bits.
- The packet header does not include Authentication and AuType fields.
- The interface ID replaces the address information in the hello packet. The Interface ID becomes the network LSA link-state ID, if the router becomes the designated router on the link.
- Router-bit (R-bit) and V6-bit in the options field process router LSAs during Shortest Path First (SPF) calculation. R-bits and V6-bits determine participation in topology distribution. The V6-bit specializes the R-bit. If the V6-bit is clear, the OSPF speaker can participate in the OSPF topology distribution without forwarding IPv6 datagrams. If the R-bit is set and the V6-bit is clear, the OSPF speaker still does not forward IPv6 datagrams, but it can forward IPv4 datagrams.
- The packet header includes the instance ID, which allows multiple OSPF protocol instances on the same link.

R-bit

Unlike OSPF for IPv4, OSPFv3 for IPv6 supports the R-bit. The R-bit indicates whether the originating node is an active router. If the R-bit is cleared, routes that transit the advertising node cannot be calculated.

For example, if a multi-homed host participates in routing without forwarding non-locally-addressed packets, the R-bit is cleared.

An IPv6-enabled switch can continue to operate as an OSPFv3 neighbor even if you disable IPv6 forwarding on the switch. This behavior differs from IPv4 OSPF, in which the switch drops a neighbor, if IP forwarding on the neighbor is disabled.

LSAs

OSPFv3 includes link LSAs and Intra-Area-Prefix LSAs.

Link LSA

The link LSA uses link flooding scope, not flooded beyond the associated link.

Link LSAs have three purposes:

- to provide the link-local address of the router to all other nodes on the link
- to provide the list of IPv6 prefixes associated with the link
- to allow the router to associate options bits with the network LSA for the link

Intra-Area-Prefix LSA

The Intra-Area-Prefix-LSA carries all IPv6 prefix information. In IPv4, this information is in router LSAs and network LSAs.

Unknown LSA types

In OSPFv3, unknown LSA types are either stored and flooded as though understood or given link flooding scope. Specific behavior is coded in the LS type field of the header.

Link LSA Suppression

To decrease unnecessary link LSA generation and flooding for non-broadcast and non-NBMA interfaces, the Link LSA Suppression interface configuration parameter has been added in RFC 5340. If Link LSA Suppression is configured for an interface, and the interface type is not broadcast or NBMA, the originated link LSA may be suppressed. Link LSA suppression is disabled, by default. For more information on configuration see, Configuring OSPF on a port or VLAN on page 175.

Stub area

OSPFv3 retains the concept of stub areas, which minimize link-state databases and routing table sizes.

IPv6 stub areas carry only router LSAs, network LSAs, Inter-Area-Prefix-LSAs, link LSAs, and Intra-Area-Prefix-LSAs.

Unlike IPv4, IPv6 can store LSAs with unrecognized link-state (LS) types or flood them as though they are understood. Rules applied to the stub area prevent the excessive growth of the link-state database. An LSA with an unrecognized link state can be flooded only if the LSA uses area- or link-flooding scope, and the LSA U-bit is 1 throughout stub and NSSA areas.

Stub area

OSPFv3 retains the concept of stub areas, which minimize link-state databases and routing table sizes.

IPv6 stub areas carry only router LSAs, network LSAs, Inter-Area-Prefix-LSAs, link LSAs, and Intra-Area-Prefix-LSAs.

Unlike IPv4, IPv6 can store LSAs with unrecognized link-state (LS) types or flood them as though they are understood. Rules applied to the stub area prevent the excessive growth of the link-state database. An LSA with an unrecognized link state can be flooded only if the LSA uses area- or link-flooding scope, and the LSA U-bit is 0.

Deprecation of MOSPF for IPV6

OSPFv3 in RFC 5340 deprecates Multicast Extensions to OSPF (MOSPF) support, and its attendant protocol fields.

NSSA Specification

RFC 2740 partially specifies this protocol feature, the level of specification was insufficient to implement it. However, RFC 5340 includes an NSSA specification unique to OSPFv3. This specification coupled with NSSA provide sufficient specification for implementation. Current Infinity IPv6 OSPF has full support for NSSA feature and is consistent with the additional specifications in RFC 5340.

Stub Area Unkown LSA Flooding Restriction Deprecated

In RFC 2740, flooding of unknown LSA was restricted within stub and NSSA areas. Following were the restrictions:

- Unlike IPv4, in IPv6 you can label unrecognized LS types as "Store and flood the LSA, as if type understood". Uncontrolled introduction of such LSAs could cause a stub area's link-state database to grow larger than its component router's capacities
- To guard the above situation, the following rule regarding stub areas has been established:

An LSA whose LS type is unrecognized can be flooded only into a stub area, if both the LSAs have area or link-local flooding scope, and the LSA has U-bit set to 0

Now the above restrictions have been deprecated. OSPFv3 routers flood link and area scope LSAs whose LS type is unrecognized and U-bit is set to 1 throughout stub and NSSA areas. The only backward compatibility issue is that the OSPFv3 routers still supporting the restrictions may not propagate newly defined LSA types.

LSA Options and Prefix Options Updates

The LSA Options and Prefix Options fields have been updated to reflect recent protocol additions. Specifically, bits related to MOSPF have been deprecated, Options field bits common with OSPFv2 have been reserved, and the DN-bit has been added to the prefix- options.

IPv6 Site-Local Address

All references to IPv6 site-local addresses have been removed in RFC 5340. Infinity IPv6 OSPF does not contain any reference to IPv6 site-local addresses and is already compliant with RFC 5340 for this.

IPsec support with OSPFv3

You can use Internet Protocol Security (IPsec) with OSPFv3 virtual link for the security protection of communication between the end points. You can also use IPsec with OSPFv3 on a brouter port or VLAN interface, for example, if you want to encrypt OSPFv3 control traffic on a broadcast network.

OSPF virtual link provides connectivity to the OSPF backbone area for redundancy or to provide a virtual link if a physical connection is not possible.

Because the device does not know the IPv6 addresses of the OSPFv3 virtual link end points at the time of configuration, you cannot manually configure the security policy ahead of time. The system must self-manage its security policy dynamically. The device also dynamically manages the IPsec enable flag, which the virtual link uses on a Layer 2 interface, either a VLAN or brouter port interface.

The following events can trigger an IPsec policy activation:

- 1. An OSPFv3 routing module detects the establishment of a virtual link.
- 2. IPsec is enabled on the already established virtual link.

On the other hand, the following two events can dynamically trigger an IPsec policy deactivation:

- 1. The virtual link is turn down.
- 2. IPsec is disabled on the virtual link.

IPsec policies can also change dynamically if a neighbor address or a local address changes.

You can enable IPsec support for IPv6 OSPF virtual link at the system level through CLI. You must disable IPsec before you can perform virtual link policy configuration changes.

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You can configure the direction you want IPsec to protect, either, ingress, egress, or both. In addition, you can permit or drop communication for the OSPF virtual link.

You can also use IPsec with OSPFv3 on a brouter port or VLAN interface. For a full configuration example and more information on IPsec, see Configuring Security for VOSS.

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.

ECMP with OSPFv3

The ECMP feature supports and complements OSPFv3 protocol.

With Equal Cost Multipath (ECMP), you can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By

maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.



Note:

To add OSPFv3 equal cost paths in the routing table, you must first enable IPv6 ECMP feature

For scaling information on the ECMP paths supported per destination prefix, see Release Notes for VOSS.

OSPFv3 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 OSPFv3 routes exist in a router and they must travel through a BGP network, then configure redistribution of OSPFv3 routes through BGP. This configuration sends OSPFv3 routes to a router that uses BGP.

You can redistribute routes on a global basis between protocols on a single VRF instance (intraVRF).

Use the ipv6 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 a redistribute policy (OSPFv3) 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.

OSPFv3 Configuration using CLI

Use the procedures in this section to configure OSPFv3 using CLI.

Configuring OSPF globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

enable
configure terminal
Optional: router vrf WORD<1-16>

2. Enable OSPFv3 for IPv6:

router ospf ipv6-enable

The default is disabled.

3. Log on to OSPF Router Configuration mode:

router ospf

4. Specify the router ID:

ipv6 router-id {A.B.C.D}

5. Optionally, make the router an autonomous system (AS) boundary router (BR):

ipv6 as-boundary-router enable

Enable the ASBR if the router attaches at the edge of the OSPF network, and has one or more interfaces that run an interdomain routing protocol. The default is disabled.

Example

Enable OSPFv3 for IPv6:

Switch:1(config) #router ospf ipv6-enable

Log on to OSPF Router Configuration mode:

Switch:1(config) #router ospf

Specify the router ID:

Switch:1(config-ospf)#ipv6 router-id 1.1.1.1

Variable definitions

Use the data in the following table to use the ipv6 router-id command.

Variable	Value
{A.B.C.D}	Specifies a 32–bit integer that identifies the router in the autonomous system. This value must be unique. The default value will be one of the IPv4 interface addresses.

Creating an OSPF area

Create an area to subdivide the autonomous system (AS) into areas that group contiguous networks, routers that connect to these networks, and attached hosts. 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.

About this task

A stub area does not receive advertisements for external routes, which reduces the size of the link-state database (LSDB). A stub area uses only one area border router (ABR). Any packets destined for outside the area are routed to the area border exit point, examined by the ABR, and forwarded to a destination.

A not so stubby area (NSSA) prevents the flooding of AS-External link-state advertisements into the area by replacing them with a default route. NSSAs also import small stub (non- OSPF) routing domains into OSPF.

Before you begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Specify the area ID:

```
ipv6 area {A.B.C.D}
```

- 3. Configure optional area parameters if the default values do not meet your requirements:
 - a. Configure the area type if you need a stub or NSSA area:

```
ipv6 area {A.B.C.D} type <nssa|stub>
```

By default, the area is a normal area; neither a stub nor NSSA area.

b. Configure the default cost:

```
ipv6 area {A.B.C.D} default-cost <0-16777215>
```

You do not need to configure this parameter if the area is a normal area.

c. Configure the area support for importing advertisements:

```
ipv6 area {A.B.C.D} import <external|noexternal|nssa>
```

The default is external.

d. Disable the importation of summary advertisements into a stub area:

no ipv6 area {A.B.C.D} import-summaries enable

The default is enabled.

e. Configure translation of Type 7 LSAs into Type 5 LSAs:

ipv6 area {A.B.C.D} translator-role <1-2>

The default value is 2-candidate.

Example

Specify the area ID:

Switch:1(config-ospf)#ipv6 area 0.0.0.1

Variable definitions

Use the data in the following table to use the ipv6 area command.

Variable	Value
{A.B.C.D}	Specifies a 32-bit integer to uniquely identify an area. Use 0.0.0.0 for the OSPFv3 backbone.
default-cost <0-16777215>	Configures the metric value advertised for the default route to stub and NSSA areas.
import <external noexternal nssa></external noexternal nssa>	Configures area support for importing AS-external LSAs:
	external—normal area
	noexternal—stub area
	nssa—not-so-stubby area
	AS-scope LSAs are not imported into stub areas or NSSAs. NSSAs import AS-External data at Type 7 LSAs, which use area scope. The default is external.
import-summaries enable	Controls the import of inter-area LSAs into a stub area. If you disable this parameter, the router does not originate nor propagate inter-area LSAs into the stub area. If you enable this parameter (the default), the router both summarizes and propagates interarea LSAs.
<nssa stub></nssa stub>	Configures the type of area. By default, the area is neither a stub area or an NSSA.
translator-role <1–2>	Indicates if the NSSA border router can perform NSSA translation of Type 7 LSAs to Type 6 LSAs. The possible values are always (1) or candidate (2). The default is candidate (2).

Creating OSPF area ranges

Create an area address range on the OSPF router to reduce the number of area border router (ABR) advertisements into other OSPF areas. An area address range is an implied contiguous range of area network addresses for which the ABR advertises a single summary route.

Before you begin

· You must create the OSPF area.

About this task

If you create two ranges, and one range is a subset of the other, the router uses the most specific match.

Before you begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create an area range:

ipv6 area range {A.B.C.D} WORD < 0-255 > [inter-area-prefix-link|nssa-extlink] advertise-mode <advertise|not-advertise> [advertise-metric < 0-65535 >]

Example

Create an area range:

Switch:1(config-ospf) #ipv6 area range 0.0.0.1 3000::0/16 advertise-mode advertise

Variable definitions

Use the data in the following table to use the ipv6 area range command.

Variable	Value
{A.B.C.D}	Specifies the area in which the address aggregate exists. Use dotted decimal notation to specify the area name.
advertise-metric <0-65535>	Specifies a cost value to advertise for the OSPF area range. This value applies to summary LSAs (Type 3).

Variable	Value
	If the value is 0, OSPF uses the cost to the farthest point in the network that is summarized.
advertise-mode <advertise not-advertise></advertise not-advertise>	Specifies the advertisement mode for prefixes in the range. advertise advertises the aggregate summary LSA with the same link-state ID. not-advertise does not advertise networks that fall within the range. The default is advertise.
<inter-area-prefix-link nssa-extlink></inter-area-prefix-link nssa-extlink>	Specifies the area LSDB type to which the address aggregate applies. inter-area-prefix-link generates an aggregated summary. nssa-extlink generates an NSSA link summary.
WORD<0-255>	Specifies the IPv6 address and prefix.

Creating an OSPF virtual link

Create a virtual link if the switch does not connect directly to the backbone. The switch can create automatic virtual links or you can perform this procedure to create virtual links manually. Manual virtual links conserve resources and provide specific control over virtual link placement in your OSPF configuration.

Before you begin

The router must be an ABR to create a virtual router interface.

About this task

Virtual linking is similar to backup redundancy. The switch creates a virtual link for vital traffic paths in your OSPF configuration if traffic is interrupted, such as when an interface cable that provides a connection to the backbone (either directly or indirectly) is disconnected from the switch. Automatic virtual linking ensures that a link is created by using another switch.

OSPF routes cannot be learned through an ABR unless it connects to the backbone directly or through a virtual link.

Before you begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create a virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D}
```

- 3. Configure optional parameters for the virtual link if the default values do not meet your requirements:
 - a. Configure the router dead interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} dead-interval
<1-65535>
```

The default is 60 seconds.

b. Configure the hello interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} hello-interval
<1-65535>
```

The default is 10 seconds.

c. Configure the retransmit interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} retransmit-interval
<1-1800>
```

The default is 5 seconds.

d. Configure the transit delay:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} transit-delay <1-1800>
```

The default is 1 second.

Example

Create a virtual link:

```
Switch:1(config-ospf) #ipv6 area virtual-link 0.0.0.1 2.2.2.2
```

Configure optional parameters for a virtual link:

```
Switch:1(config-ospf) #ipv6 area virtual-link 0.0.0.1 4.4.4.4 dead-interval 90 retransmit-interval 10
```

Variable definitions

Use the data in the following table to use the ipv6 area virtual-link command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the ID for the transit area that the virtual link traverses and the router ID of the virtual neighbor. Do not use 0.0.0.0 for the transit area.
dead-interval <1-65535>	Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets. Configure this value as a multiple of the hello interval. You must configure the same value on the virtual neighbor. The default is 60 seconds.

Variable	Value
hello-interval <1-65535>	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor. The default is 10 seconds.
retransmit-interval <1-1800>	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets. The default is 5 seconds.
transit-delay <1-1800>	Specifies the estimated number of seconds to transmit a link-state update packet over this interface. The default is 1 second.

Configuring IPsec for the OSPF virtual link

Use the following procedure to configure and enable IPsec for the OSPF virtual link.

IPsec is disabled by default.

Before you begin

- Configure the OSPF virtual link.
- Create the IPsec security association. For more information on configuration of IPsec security associations and IPsec policies, and how to enable policies on an interface, see <u>Configuring</u> <u>Security for VOSS</u>.

About this task

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You must disable IPsec before you can perform virtual link policy configuration changes.

For configuration examples of IPsec used with OSPFv3 virtual link, see <u>Configuring Security for VOSS</u>.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create the IPsec policy under the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec
```

3. Configure the action of the IPsec policy under the OSPF virtual link:

ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec action <drop|
permit>

4. Configure the direction of the IPsec policy under the OSPF virtual link:

ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec direction <both|in|
out>

5. Link the security association to the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec security-association WORD<0-32>
```

6. Enable the IPsec policy created under the OSPF virtual link:

ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec enable

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#(config)router ospf
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec action permit
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec direction both
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec security-association
test1
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec enable
```

Variable definitions

Use the data in the following table to use the ipv6 area virtual link {A.B.C.D} {A.B.C.D} ipsec Command.

Variable	Value
{A.B.C.D}{A.B.C.D}	The first IP address specifies the area IP address, and the second IP address specifies the virtual-link IP address.
action <drop permit></drop permit>	Configures the action of the IPsec policy under the OSPF virtual tunnel to one of the following:
	drop—Drops the IP packets.
	permit—Permits the IP packets.
	The default is permit.
direction <both in out></both in out>	Specifies the direction you want to protect with IPsec:
	in—Specifies ingress traffic.
	out—Specifies egress traffic.
	both—Specifies both ingress and egress traffic.
	The default is both.
enable	Enables the IPsec policy under the OSPF virtual link.

Variable	Value
security-association WORD<0-32>	Links the security association to the OSPF virtual link.

Configuring OSPF default metrics

Before you begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

About this task

Use the following procedure to configure global OSPF default metrics.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Configure OSPF default-cost:

```
ipv6 default-cost {ethernet|fast-ethernet|forty-qiq-ethernet|
hundred-gig-ethernet|gig-ethernet|ten-gig-ethernet|twentyfive-gig-
ethernet|vlan} <1-65535>
```



Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Example

Configure IPv6 default cost metric for Ethernet to 100, for fast Ethernet to 20, for gig-ethernet, twentyfive-gig-ethernet, forty-gig-Ethernet, and hundred-gig-ethernet to 2, and VLAN to 1.

```
Switch:1>enable
Switch: 1#configure terminal
Switch:1(config) #router ospf
Switch:1(config-ospf) #ipv6 default-cost ethernet 100
Switch:1(config-ospf) #ipv6 default-cost fast-ethernet 20
Switch:1(config-ospf) #ipv6 default-cost gig-ethernet 2
Switch:1(config-ospf) #ipv6 default-cost ten-gig-ethernet 2
Switch:1(config-ospf) #ipv6 default-cost Forty-gig-ethernet 2
Switch:1(config-ospf) #ipv6 default-cost twentyfive-gig-ethernet 2
Switch:1(config-ospf) #ipv6 default-cost hundred-gig-ethernet 2
Switch:1(config-ospf) #ipv6 default-cost vlan 1
```

Variable definitions

Use the data in the following table to use the ipv6 default-cost command.

Note:

Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Variable	Value
ethernet <1-65535>	Configures the IPv6 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 IPv6 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 IPv6 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 10 Mb/s Forty-Gigabit- Ethernet (default is 1).
gigabit-ethernet <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
	gigabit-ethernet is for 10 Mb/s Gigabit-Ethernet (default is 1).
hundred-gig-ethernet <1-65535>	Configures the IPv6 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 IPv6 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 Mb/s Ten-Gigabit-Ethernet (default is 1).
twentyfive-gig-ethernet <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
	On a channelized 100 Gbps port, the default-cost for each 25 Gbps channel is 1.
vlan <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
	vlan is for Vlan interfaces (default is 10).

Configuring OSPF on a port or VLAN

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface.

Before you begin

· The IPv6 interface must exist.

Procedure

1. Enter Interface Configuration mode:

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

Note:

If the 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 interface:

```
ipv6 ospf area {A.B.C.D}
```

3. Enable OSPFv3 on the interface:

```
ipv6 ospf enable
```

The default is disabled.

- 4. Configure optional parameters to meet your requirements:
 - a. Configure the interface metric:

```
ipv6 ospf cost <0-65535>
```

The default for a brouter port or VLAN is 1.



If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.

b. Configure the router dead interval:

```
ipv6 ospf dead-interval <1-65535>
```

The default is 40 seconds.

c. Configure the hello interval:

```
ipv6 ospf hello-interval <1-65535>
```

The default is 10 seconds.

d. Configure the link LSA suppression:

ipv6 ospf link-lsa-suppression



Note:

Before configuring Link LSA suppression for OSPF, configure Link LSA suppression for OSPF area for point to point (p2p) or point to multipoint interfaces (p2mp), otherwise it defaults to a broadcast interface type where you cannot use Link LSA suppression.

e. Configure the poll interval:

ipv6 ospf poll-interval <0-65535>

The default is 120 seconds.

f. Configure the interface priority:

ipv6 ospf priority <0-255>

The default is 1.

g. Configure the retransmit interval:

ipv6 ospf retransmit-interval <1-1800>

The default is 5 seconds.

h. Configure the transit delay:

ipv6 ospf transit-delay <1-1800>

The default is 1 second.

Example

Create an OSPF area on the interface:

Switch:1(config-if)#ipv6 ospf area 0.0.0.0

Enable OSPFv3 on the interface:

Switch:1(config-if) #ipv6 ospf enable

Variable definitions

Use the data in the following table to use the ipv6 ospf command.

Variable	Value
area {A.B.C.D}	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
cost <0-65535>	Specifies the cost for the interface.
	The default for a brouter port or VLAN is 1.
dead-interval <1-65535>	Specifies the number of seconds after which the neighbor declares the router down, if it does not

Variable	Value
	receive hello packets. Configure this value as a multiple of the hello interval. You must configure the same value on the virtual neighbor.
	The default is 40 seconds.
enable	Specifies the administrative status for the OSPFv3 interface.
	If you enable the status, it is advertised as an interal route to some areas.
	If you disable the status, the interface is external to OSPFv3.
	The default is disabled.
hello-interval <1-65535>	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor.
	The default is 10 seconds.
link-lsa-suppression	Configures link LSA suppression on the specified port or VLAN. It is only used for point to point or point to multipoint interfaces. By default, it is disabled.
poll-interval <0-65535>	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor.
	The default is 120.
priority <0-255>	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router.
	The default is 1.
retransmit-interval <1-1800>	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets. The default is 5 seconds.
transit-delay <1-1800>	Specifies the estimated number of seconds to transmit a link-state update packet over this interface.
	The default is 1 second.

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring OSPF on a tunnel

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface.

Before you begin

• The IPv6 interface must exist.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create an OSPF area on the interface:

```
ipv6 tunnel <1-2000> area {A.B.C.D}
```

3. Enable OSPFv3 on the interface:

```
ipv6 tunnel <1-2000> enable
```

- 4. Configure optional parameters to meet your requirements:
 - a. Configure the router dead interval:

```
ipv6 tunnel <1-2000> dead-interval <1-65535>
```

The default is 40 seconds.

b. Configure the hello interval:

```
ipv6 tunnel <1-2000> hello-interval <1-65535>
```

The default is 10 seconds.

c. Configure the interface metric:

```
ipv6 tunnel <1-2000> metric <0-65535>
```

d. Configure the poll interval:

```
ipv6 tunnel <1-2000> poll-interval <0-65535>
```

The default is 120 seconds.

e. Configure the interface priority:

```
ipv6 tunnel <1-2000> priority <0-255>
```

The default is 1.

f. Configure the retransmit interval:

```
ipv6 tunnel <1-2000> retransmit-interval <1-1800>
```

The default is 5 seconds.

g. Configure the transit delay:

```
ipv6 tunnel <1-2000> transit-delay <1-1800>
```

The default is 1 second.

Example

Create an OSPF area on the interface:

```
Switch:1(config-if) #ipv6 tunnel 4 area 0.0.0.0
```

Enable OSPFv3 on the interface:

Switch:1(config-if) #ipv6 tunnel 4 enable

Variable definitions

Use the data in the following table to use the ipv6 tunnel command.

Variable	Value
<1–2000>	Specifies the tunnel ID.
area {A.B.C.D}	Specifies the area ID to which the IPv6 interface connects.
	Use 0.0.0.0 for the OSPFv3 backbone.
dead-interval <1-65535>	Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets.
	Configure this value as a multiple of the hello interval.

Variable	Value
	Tip:
	You must configure the same value on the virtual neighbor.
	The default is 40 seconds.
enable	Specifies the administrative status for the OSPFv3 interface.
	If you enable the status, it is advertised as an internal route to some areas.
	If you disable the status, the interface is external to OSPFv3.
	The default is enabled.
hello-interval <1-65535>	Specifies the number of seconds between hello packets that the router sends on this interface.
	Tip:
	You must configure the same value on the virtual neighbor.
	The default is 10 seconds.
metric <0-65535>	Specifies the cost for the interface.
	The default for a tunnel is 100.
poll-interval <0-65535>	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor.
	The default is 120.
priority <0-255>	Specifies the priority of this interface.
	Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the likelihood that the router becomes the designated router.
	A value of zero (0) indicates the router cannot become the designated router for the network.
	If more than one router uses the same priority value, the system uses the router ID to determine which router becomes the designated router.
	The default is 1.
retransmit-interval <1-1800>	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface.

Variable	Value
	The retransmit-interval value also applies to the retransmissions of database description and link-state request packets.
	The default is 5 seconds.
transit-delay <1-1800>	Specifies the estimated number of seconds required to transmit a link-state update packet over this interface.
	The default is 1 second.

Viewing OSPFv3 Information

View information about OSPF to view the current configuration.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View OSPF global information:

```
show ipv6 ospf [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

3. View OSPF areas:

```
show ipv6 ospf area [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

4. View OSPF interface information

5. View OSPF interface timers:

```
show ipv6 ospf int-timers [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

6. View the link-state database (LSDB) table:

```
show ipv6 ospf lsdb [adv-rtr <A.B.C.D>] [area <A.B.C.D>] [interface gigabitEthernet {slot/port[sub-port]}|vlan <1-4059>] [lsa-type <1-11>] [lsid <0-4294967295>] [scope <1-3>] [tunnel <1-2000>] [detail] [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

7. View OSPF neighbors to see routers with interfaces to a common network, including neighbors on the virtual link to the OSPF backbone:

```
show ipv6 ospf neighbor [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

Example



```
OSPFv3 Global Information - GlobalRouter
: 170.76.84.0
     router-id
                          : DISABLE
     admin-state
                          : 3
     version
     area-bdr-rtr-state
                          : FALSE
     as-bdr-rtr-state
                          : FALSE
     helper-mode
                          : ENABLED
     as-scope-lsa-count
                          : 0
     lsa-checksum
                          : 0
     originate-new-lsas
                          : 0
     rx-new-lsas
                          : 0
     ext-lsa-count
                          : 0
         default-metric :
                ethernet - 100
             fast-ethernet - 10
              gig-ethernet - 1
           ten-gig-ethernet - 1
          forty-gig-ethernet - 1
                  vlan - 10
Switch:1>show ipv6 ospf area
______
                   OSPF Area - GlobalRouter
AREA ID STUB AREA NSSA IMPORT SUM TRANS ROLE TRANS STATE
0.0.0.0 false false true always disabled
STUB_METRIC STUB_METRIC_TYPE SPF_RUNS BDR_RTR_CNT ASBDR_RTR_CNT LSA_CNT LSACK_SUM
10 ospfV3Metric 0 0 0 0 0
Switch: 1#show ipv6 ospf interface
Total ospf areas: 1
Total ospf interfaces: 2
______
                    OSPF Interface - GlobalRouter
______
IFINDX(VID/BRT) AREAID ADM IFSTATE METRIC PRI DR/BDR IFTYPE
331 (5/12 ) 0.0.0.0 ena DOWN 1 1 0.0.0.0 PT-PT 0.0.0.0 lnklsaSup 2078 (30 ) 0.0.0.0 ena DR 1 1 197.146.128.0 BROADCAST 0.0.0.0
2 out of 2 Total Num of ospf interfaces displayed
Total ospf virtual interfaces: 0
______
              OSPF Virtual Interface - GlobalRouter
______
AREAID NBRIPADDR STATE
0 out of 0 Total Num of ospf virtual interfaces displayed
Switch: 1#show ipv6 ospf int-timers
```

OSPF Interface Timers - GlobalRouter									
	X (VID/BR		AREAID		TRANSIT	RETRANS INTERVAL	HELLO	DEAD	POLL INTERVAL
2059 2060	(11 (12)	0.0.0.0		1	5 5	10 10		120 120
=====		==:	OSPF	===: Vir	tual Inte	rface Time	ers - Glob	oalRouter	
AREAI	D 		NBRIPADDR				HELLO INTERVAL	DEAD INTERVAL	
Switc	h:1#show		pv6 ospf neig						
				OS	PF Neighb	or - Globa	alRouter		
			NBRROUTERID		NBRIPADD			 ГАТЕ 	TTL
331 (10/19) 97.146.128.0 fe80:0:0:0:2ef4:c5ff:fe92:8a00 Restart 120									
1 out of 1 Total Num of Neighbor Entries displayed.									
OSPF Virtual Neighbor - GlobalRouter									
NBRAREAID NBRROUTERID VIRTINTFID NBRIPV6ADDR STATE									
0 out of 0 Total Num of Virtual Neighbor Entries displayed.									
OSPF NBMA Neighbor - GlobalRouter									
	FACE NBR				ADDR				======= STATE
0 out of 0 Total Num of NBMA Neighbor Entries displayed.									

Variable Definitions

Use the data in the following table to use the **show ipv6** ospf commands.

Variable	Value
adv-rtr <a.b.c.d></a.b.c.d>	Shows information for the specified advertising router.
area <a.b.c.d></a.b.c.d>	Shows information for the specified area.
detail	Shows information beyond the basic information.

Variable	Value
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
Isa-type <1-11>	Shows information for the specified LSA type.
Isid <0-4294967295>	Shows information for the specified link-state ID.
scope <1-3>	Shows information for the specified scope:
	 link-scope LSAs-View the link-scope LSDB to view the LSAs that are not flooded beyond the local link.
	 area-scope LSAs-View the area-scope LSDB to see the LSAs that are flooded in a single OSPF area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter- Area-Router LSAs, and Intra-Area-Prefix-LSAs.
	 AS-scope LSAs-View the AS-scope LSDB to see the LSAs that are flooded through the routing domain. The AS scope is used for ASexternal- LSAs.
tunnel <1-2000>	Specifies the ID number of the tunnel.
vlan <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.
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies VRF IDs.

Viewing OSPFv3 Default Cost Information

About this task

Use the following procedure to view the OSPF default cost information, to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the OSPF cost information:

show ipv6 ospf default-cost [vrf WORD < 1-16 >] [vrfids WORD < 0-512 >]

Example

```
Switch:1#show ipv6 ospf default-cost

IPv6 OSPF Default Metric - GlobalRouter

10MbpsPortDefaultMetric: 100
100MbpsPortDefaultMetric: 1
1000MbpsPortDefaultMetric: 1
10000MbpsPortDefaultMetric: 1
25000MbpsPortDefaultMetric: 1
40000MbpsPortDefaultMetric: 1
10000MbpsPortDefaultMetric: 1
VlanDefaultMetric: 1
VlanDefaultMetric: 1
```

Adding an NBMA neighbor

Add an NBMA neighbor for each interface that is eligible to become the DR.

An NMBA interface with a positive nonzero router priority is eligible to become the DR for the NBMA network and is configured with the identification of all attached routers, IPv6 addresses, and router priorities.

Before you begin

- Identify the following information:
 - specific interfaces to include in the NBMA network
 - the IPv6 address for each interface
 - the router priority for each interface
 - the hello interval for the network
 - the router dead interval for the network
 - the poll interval for the network

About this task

In contrast to a broadcast network where switches multicast (send to AllSPFRouters and AllDRouters) certain OSPF protocol packets, switches replicate and send NBMA packets to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination addresses AllSPFRouters and AllDRouters. Because the NBMA network does not broadcast, you must manually configure a list of neighbors and priorities for all routers in the network that can become the DR. Potential DRs use a positive nonzero router priority.

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

2. Create a new NBMA neighbor:

```
ipv6 ospf nbma-nbr WORD<0-43> <0-255>
```

3. Change the priority of an existing NBMA neighbor:

```
ipv6 ospf nbma-nbr WORD<0-43> priority <0-255>
```

Example

Create an NBMA neighbor that will not become the DR:

Switch:1(config-if) #ipv6 ospf nbma-nbr fe80:0:0:8217:7dff:fe76:8a03 0

Variable definitions

Use the data in the following table to use the ipv6 ospf nbma-nbr command.

Variable	Value
priority <0-255>	Specifies the priority to use for this neighbor in the designated router election process. A value of 0 indicates the neighbor cannot become the designated router. The higher the priority value, the higher chance the switch will win the election process. The default is 1.
WORD<0-43>	Specifies the IPv6 address of the neighbor.

Configuring link LSA suppression

About this task

Use the following procedure to configure link LSA suppression on a port or a VLAN, to decrease unnecessary link LSA generation and flooding for non-broadcast and non-NBMA interface.

Procedure

1. Enter Interface Configuration mode:

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

■ Note:

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

2. Enter the following command:

```
ipv6 ospf area {A.B.C.D} network {p2p | p2mp} link-lsa-suppression
```

Example

Variable definitions

Following table describes the variables to the ipv6 ospf area {A.B.C.D} network p2p link-lsa-suppression command.

Variable	Description
area {A.B.C.D}	Create an IPv6 OSPF area.
network	Sets the type of interface.
[eth NBMA p2mp p2p passive]	Specifies the type of interface.
link-Isa-suppression	Enables link LSA suppression.

Configuring Route Redistribution to OSPFv3 in GRT mode

Configure a redistribute entry to announce certain routes into the OSPFv3 domain, including static routes, direct routes, RIPng, OSPF routes, IS-IS routes, or Border Gateway Protocol (BGP) routes. Optionally, use a route policy to control the redistribution of routes.

Note:

RIPng is not virtualized, therefore RIPng redistribution in OSPFv3 works only in GRT.

Before you begin

- Enable OSPFv3 globally.
- Ensure that a route policy exists.
- Ensure that you set OSPFv3 as the boundary router.

Note:

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

- For a redistribute policy (OSPFv3) 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.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

Procedure

1. Enter OSPF Router Configuration mode:

enable

configure terminal
router ospf

2. Create the redistribution instance:

ipv6 redistribute <bgp|direct|isis|rip|static>



The switch loads the existing configuration when you upgrade to the current release. Once you have the configuration file saved using the current release, only the new configuration will be loaded.

3. Apply a route policy if required:

ipv6 redistribute

bgp|direct|isis|rip|static> route-map WORD<0-64>



No inter-vrf Route Redistribution is supported for IPv6.

- 4. Configure other parameters, as required.
- 5. Enable the redistribution.

ipv6 redistribute <bgp|direct|isis|rip|static> enable

6. Ensure that the configuration is correct:

show ipv6 ospf redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]

7. Exit to Global Configuration mode:

exit

8. Apply the redistribution.

ipv6 ospf apply redistribute <bgp|direct|isis|rip|static> [vrf
WORD<1-16>]

Changes do not take effect until you apply them.

- 9. View all routes that are redistributed into OSPFv3:
 - a. View the routes that are redistributed from the GRT to OSPFv3:

```
show ipv6 ospf redistribute
```

b. View the routes that are redistributed to OSPFv3 for a specific VRF instance:

```
show ipv6 ospf redistribute [vrf WORD<1-64>] [vrfids WORD<0-512>]
```

Example

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)#ipv6 redistribute static
WARNING: Routes will not be injected until apply command is issued after enable command
Switch:1(config-ospf) #ipv6 redistribute static route-map policy1
Switch:1(config-ospf) #ipv6 redistribute static enable
Switch:1(config-ospf)#exit
Switch:1(config) #ipv6 ospf apply redistribute static
Switch:1(config) #show ipv6 ospf redistribute
                 OSPFv3 Redistribute List - GlobalRouter
______
SRC MET MTYPE
                    ENABLE
                              RPOLICY
STAT 0 type2 TRUE policy1
```

Variable Definitions

Use the data in the following table to use the ipv6 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></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.
<bgp direct isis rip static></bgp direct isis rip static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, or static.

Use the data in the following table to use the ipv6 ospf apply redistribute command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance.
 static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, or static.

Viewing the Status of OSPFv3 Redistribution

View the status of OSPFv3 route redistribution to verify the current configuration. You can redistribute directly connected routes and IPv6 static routes into OSPFv3.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the current configuration:

show ipv6 ospf redistribute [vrf WORD <1-16>] [vrfids WORD <0-512>]

Example

```
Switch:1#show ipv6 ospf redistribute

OSPFv3 Redistribute List - GlobalRouter

SOURCE MET MTYPE ENABLE RPOLICY

Static 0 external TRUE
```

Variable Definitions

Use the data in the following table to use the show ipv6 ospf redistribute command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies VRF IDs.

Disabling Helper mode for OSPFv3

Before you begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace ipv6 with ipv6 ospf.

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" commands.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Enter the following command to disable Helper mode:

```
ipv6 helper-mode-disable
```

3. Enter the following command to enable Helper mode:

```
no ipv6 helper-mode-disable
```

Or

default ipv6 helper-mode disable

Example

Disabling Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #router ospf
Switch:1(config-ospf) #ipv6 helper-mode-disable
```

Enabling Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #router ospf
Switch:1(config-ospf) #no ipv6 helper-mode-disable
```

OSPFv3 Configuration using EDM

Use the procedures in this section to configure OSPFv3 using EDM.

Configuring OSPFv3 globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Before you begin

• Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPSec on OSPFv3 virtual link interfaces

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.
- 3. Click the Globals tab.
- 4. Type the router ID, in the format of an IPv4 address.
- 5. Select enabled.
- Optionally, select ASBdrRtrStatus to make the router an AS boundary router.

Enable the ASBR if the router attaches at the edge of the OSPF network, and has one or more interfaces that run an interdomain routing protocol. The default is disabled.

7. Click Apply.

Globals field descriptions

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

Note:

Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Name	Description
Routerld	Specifies a 32–bit integer that identifies the router in the autonomous system. This value must be unique. The default value will be one of the IPv4 interface addresses.
AdminStat	Enables or disables OSPFv3 on the router. If you disable OSPFv3 globally, you disable it on all interfaces. The default is disabled.
VersionNumber	Shows the OSPF version number, which for IPv6 is version 3.
AreaBdrRtrStatus	Shows if the router is an area border router.
ASBdrRtrStatus	Configures the router as an autonomous system boundary router. The default is disabled (clear).
HelperModeDisable	Disables Graceful Restart Helper Mode feature.
AsScopeLsaCount	Shows the number of AS-external link-state advertisements in the LSDB.
AsScopeLsaCksumSum	Shows the sum of the checksums for the link-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.
OriginateNewLsas	Shows the number of new link-state advertisements. The number increases each time the router originates a new LSA.
RxNewLsas	Shows the number of new link-state advertisements received. This number does not include new instances of self-originated link-state advertisements.
ExtLsaCount	Shows the number of external (LS type 0x4005) LSAs in the LSDB.
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.

Name	Description
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	Indicates the default cost applied to VLAN interfaces. The default is 10.

Creating an OSPFv3 Area

Create an area to subdivide the autonomous system (AS) into areas that group contiguous networks, routers that connect to these networks, and attached hosts. 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.

About this task

A stub area does not receive advertisements for external routes, which reduces the size of the linkstate database (LSDB). A stub area uses only one area border router (ABR). Any packets destined for outside the area are routed to the area border exit point, examined by the ABR, and forwarded to a destination.

A not so stubby area (NSSA) prevents the flooding of AS-External link-state advertisements into the area by replacing them with a default route. NSSAs also import small stub (non- OSPF) routing domains into OSPF.

Before you begin

• Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPSec on OSPFv3 virtual link interfaces

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.
- 3. Click the Areas tab.
- 4. Click Insert.

- 5. Type the area ID.
- 6. Click Insert.

Areas field descriptions

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

Name	Description
Id	Specifies a 32-bit integer to uniquely identify an area. Use 0.0.0.0 for the OSPFv3 backbone.
ImportasExtern	Indicates the support for importing AS-external LSAs::
	importExternal—normal area
	importNoExternal—stub area
	• importNssa—not-so-stubby-area
	AS-scope LSAs are not imported into stub areas or NSSAs. NSSAs import AS-External data at Type 7 LSAs, which use area scope. importExternal is the default.
SpfRuns	Shows the number of times the intra-area route table was calculated using the LSDB of this area.
BdrRtrCount	Shows the number of reachable ABRs in this area. The value starts at zero (0). The system calculates this value in each SPF run.
AsBdrRtrCount	Shows the number of reachable ASBRs in this area. The value starts at zero (0). The system calculates this value in each SPF run.
ScopeLsaCount	Shows the number of area-scope LSAs in the LSDB for this area.
ScopeLsaCksumSum	Shows the sum of the checksums for the area-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.
Summary	Controls the import of inter-area LSAs into a stub area. If the value is noAreaSummary , the router does not originate nor propagate inter-area LSAs into the stub area. If the value is sendAreaSummary (the default), the router both summarizes and propagates inter-area LSAs.
StubMetric	Configures the metric value advertised for the default route to stub and NSSA areas.
NssaTranslatorRole	Indicates if the NSSA border router can perform NSSA translation of Type 7 LSAs to Type 6 LSAs.

Name	Description
	The possible values are always or candidate. The default is candidate.
NssaTranslatorState	Indicates if and how an NSSA border router translates Type 7 LSAs to Type 5 LSAs. The possible values are
	enabled—The border router always translates the LSAs.
	elected—A candidate border router translates the LSAs.
	disabled—A candidate border router does not translate the LSAs.
StubMetricType	Specifies the type of metric advertised as a default route. The possible values are:
	ospfv3Metric—OSPF metric
	comparableCost—external Type 1
	nonComparable—external Type 2
	The default is ospfv3Metric.

Creating OSPFv3 Area Ranges

Create an area address range on the OSPFv3 router to reduce the number of area border router (ABR) advertisements into other OSPF areas. An area address range is an implied contiguous range of area network addresses for which the ABR advertises a single summary route.

Before you begin

- · You must create the OSPF area.
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPSec on OSPFv3 virtual link interfaces

About this task

If you create two ranges, and one range is a subset of the other, the router uses the most specific match.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.

- 3. Click the Area Aggregate tab.
- 4. Click Insert.
- 5. Select the area ID.
- 6. Select the type of area.

interAreaPrefixLsa generates an aggregated summary.
nssaExternalLsa generates an NSSA link summary.

- 7. Type the prefix for the IPv6 area address.
- 8. Type the number of bits from the IPv6 address that you want to advertise.
- 9. 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 aggregate exists. Use dotted decimal notation to specify the area name.
AreaLsdbType	Specifies the area LSDB type to which the address aggregate applies. interAreaPrefixLsa generates an aggregated summary. nssaExternalLsa generates an NSSA link summary.
Prefix	Specifies the IPv6 prefix. The prefix and prefix length define the range.
PrefixLength	Specifies the length of the prefix, in bits. The prefix cannot be shorter than 3 bits. The prefix and prefix length define the range.
Effect	Specifies the advertisement mode for prefixes in the range. advertiseMatching advertises the aggregate summary LSA with the same link-state ID. doNotAdvertiseMatching does not advertise networks that fall within the range.
AdvertiseMetric	Specifies a cost value to advertise for the OSPF area range. This value applies to summary LSAs (Type 3). If the value is 0, OSPF uses the cost to the farthest point in the network that is summarized.

Creating an OSPFv3 Virtual Link

Create a virtual link if the switch does not connect directly to the backbone. The switch can create automatic virtual links or you can perform this procedure to create virtual links manually. Manual

virtual links conserve resources and provide specific control over virtual link placement in your OSPFv3 configuration.

Before you begin

- The router must be an ABR to create a virtual router interface.
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.

Note:

Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPSec on OSPFv3 virtual link interfaces

About this task

Virtual linking is similar to backup redundancy. The switch creates a virtual link for vital traffic paths in your OSPFv3 configuration if traffic is interrupted, such as when an interface cable that provides a connection to the backbone (either directly or indirectly) is disconnected from the switch. Automatic virtual linking ensures that a link is created by using another switch.

OSPF routes cannot be learned through an ABR unless it connects to the backbone directly or through a virtual link.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.
- 3. Click the Virtual If tab.
- 4. Click Insert.
- 5. Specify the ID for the transit area.

The transit area is the common area between two ABRs.

6. Specify the router ID for the virtual neighbor.

The neighbor ID is the IP router ID of the ABR through which the other ABR must route traffic destined for the backbone.

- 7. Click Insert.
- 8. Click **Refresh** to verify that the virtual link is active.

If the state is point-to-point, the virtual link is active. If the state is down, the virtual link is configured incorrectly.

Virtual If field descriptions

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

Name	Description
Areald	Specifies the ID for the transit area that the virtual link traverses. Do not use 0.0.0.0.
Neighbor	Specifies the router ID of the virtual neighbor.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state update packet over this interface. The default is 1 second.
RetransInterval	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets.
	The default is 5 seconds.
HelloInterval	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor. The default is 10 seconds.
RtrDeadInterval	Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets. Configure this value as a multiple of the hello interval.
	You must configure the same value on the virtual neighbor. The default is 60 seconds.
State	Shows the state of the virtual interface: either down or pointToPoint.
Events	Shows the number of state changes or error events on the virtual link.
LinkScopeLsaCount	Shows the number of link-scope LSAs in the LSDB for the virtual link.
LinkLsaCksumSum	Shows the sum of the checksums for the link-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.

Configuring IPsec for the OSPF virtual link

Use the following procedure to configure and enable IPsec for the OSPF virtual link.

IPsec is disabled by default.

About this task

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You must disable IPsec before you can perform virtual link policy configuration changes.

Before you begin

- · Configure the OSPF virtual link.
- Create the IPsec security association.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Security > Control Panel**.
- 2. Click IPSec.
- 3. Click the OSPF Virtual Link tab.
- 4. Click Insert.
- 5. Specify the area ID.
- 6. Specify the neighbor address.
- 7. Complete the remaining optional configuration.
- 8. Click Insert.

OSPF Virtual Link field descriptions

Use the data in the following table to use the **OSPF Virtual Link** tab.

Name	Description
Areald	Identifies the OSPF virtual link area.
Neighbor	Identifies the OSPF virtual link neighbor.
SAName	Links the security association to the OSPF virtual link.
AdminStatus	Enables the policy. The default is disabled.
Action	Configures the action of the IPsec policy under the OSPF virtual tunnel to one of the following:
	permit—Permits the IP packets.
	drop—Drops the IP packets.
	The default is permit.
Direction	Specifies the direction you want to protect with IPsec:
	inBound—Specifies ingress traffic.
	outBound—Specifies egress traffic.
	bothDirections—Specifies both ingress and egress traffic.
	The default is bothDirections.

Name	Description
SrcAddress	Shows the address of the source interface to which the policy applies.
DstAddress	Shows the address of the destination interface to which the policy applies.
LinkID	Shows a unique ID for the OSPF virtual link. The default is 0.
IfIndex	Shows the interface index to which OSPF virtual link the policy applies.
OperStatus	Shows the operational status of the link, either up or down. The default is down.

Creating an OSPF interface on a brouter port

Configure the OSPF protocol on an IPv6 interface to support dynamic routing on the interface. Perform this procedure to create an OSPF interface on a brouter port.

If you want to modify existing OSPFv3 interfaces, see <u>Modifying an OSPF interface</u> on page 213. To configure OSPFv3 on an IPv6 VLAN, see <u>Creating an OSPF VLAN interface</u> on page 202.

Before you begin

• The IPv6 interface must exist.

Procedure

- 1. In the Device Physical view, select a port.
- 2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
- 3. Click IPv6.
- 4. Click the IPv6 OSPF Interface tab.
- 5. Click Insert.
- 6. Select the area ID.
- 7. Select enabled.
- 8. Click Insert.

IPv6 OSPF Interface field descriptions

Use the data in the following table to use the IPv6 OSPF Interface tab.

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.

Name	Description
Areald	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Туре	Specifies the OSPFv3 interface type as one of the following:
	broadcast
	• NBMA
	• point-to-point
	point-to-multipoint
	• passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an interal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router.
	The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelioInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this field to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this field to the same value for all routers attached to a common network. The default is 40.

Name	Description
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following:
	• down
	loopback
	waiting
	pointToPoint
	designatedRouter
	backupDesginatedRouter
	otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100.
	Note:
	If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.
BfdEnable	Enable Bidirectional Forwarding Detection (BFD) for OSPF.

Create an OSPF VLAN Interface

Configure the OSPF protocol on an IPv6 VLAN to support dynamic routing on the interface.

If you want to modify existing OSPFv3 interfaces, see Modifying an OSPF interface on page 213.

Before you begin

• The IPv6 interface must exist.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
- 2. Select VLANs.

- 3. Select the **Basic** tab.
- 4. Select a VLAN.
- 5. Select IPv6.
- 6. Select the IPv6 OSPF Interface tab.
- 7. Select Insert.
- 8. Select the area ID.
- 9. Select enabled.
- 10. Select Insert.

IPv6 OSPF Interface field descriptions

Use the data in the following table to use the IPv6 OSPF Interface tab.

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
Areald	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Туре	Specifies the OSPFv3 interface type as one of the following:
	broadcast
	• NBMA
	• point-to-point
	point-to-multipoint
	• passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an interal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router.
	The default is 1.

Name	Description
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this field to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this field to the same value for all routers attached to a common network. The default is 40.
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following:
	• down
	loopback
	waiting
	pointToPoint
	designatedRouter
	backupDesginatedRouter
	otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100.
	Note:
	If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost.

Name	Description
	The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for OSPF.

Creating an OSPF interface on a tunnel

Configure the OSPF protocol on an IPv6 interface to support dynamic routing on the interface. Perform this procedure to create an OSPF interface on a tunnel.

If you want to modify existing OSPFv3 interfaces, see <u>Modifying an OSPF interface</u> on page 213. To configure OSPFv3 on an IPv6 VLAN, see <u>Creating an OSPF VLAN interface</u> on page 202.

Before you begin

· The IPv6 interface must exist.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click Tunnel.
- 3. Click the **Tunnel Config** tab.
- 4. Select a configured tunnel.
- 5. Click IPv6 OSPF.
- 6. Click Insert.
- 7. Select the area ID.
- 8. Select enabled.
- 9. Click Insert.

OSPF Interface field descriptions

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

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
Areald	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.

Name	Description
Туре	Specifies the OSPFv3 interface type as one of the following:
	broadcast
	• NBMA
	• point-to-point
	point-to-multipoint
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an interal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router.
	The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this field to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this field to the same value for all routers attached to a common network. The default is 40.
Pollinterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.

Name	Description
State	Shows the state of the OSPFv3 interface as one of the following:
	• down
	loopback
	waiting
	pointToPoint
	designatedRouter
	backupDesginatedRouter
	otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100.
	Note:
	If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.

Viewing the AS-scope link-state database

View the AS-scope link-state database (LSDB) to see the LSAs that are flooded through the routing domain. The AS scope is used for AS external-LSAs.

Before you begin

• Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPFv3.
- 3. Click the **AS-scope LSDB** tab.

AS-scope LSDB field descriptions

Use the data in the following table to use the AS-scope LSDB tab.

Name	Description
Туре	Shows the type of the link-state advertisement. Each link state type has a separate advertisement format. AS-scope LSAs not recognized by the router may be stored in the database.
RouterId	Shows the 32 bit number that uniquely identifies the originating router in the autonomous system.
Lsid	Identifies the piece of the routing domain that is being described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Indicates the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

Viewing the area-scope LSDB

View the area-scope LSDB to see the LSAs that are flooded in a single OSPFv3 area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs.

Before you begin

• Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPFv3.
- 3. Click the **Area-scope LSDB** tab.

Area-scope LSDB field descriptions

Use the data in the following table to use the Area-scope LSDB tab.

Name	Description
Areald	Identifies the area ID from which the LSA is received. Area ID 0.0.0.0 is the OSPF backbone.

Name	Description
Туре	Identifies the type of the link-state advertisement. Each link-state type has a separate advertisement format. Area-scope LSAs unrecognized by the router are also stored in this database.
Routerid	Identifies the originating router in the autonomous system.
Lsid	Identifies the piece of the routing domain that is described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Indicates the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

Viewing the link-scope LSDB

View the link-scope LSDB to view the LSAs that are not flooded beyond the local link.

Before you begin

• Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPFv3.
- 3. Click the Link-scope LSDB tab.

Link-scope LSDB field descriptions

Use the data in the following table to use the **Link-scope LSDB** tab.

Name	Description
IfIndex	Shows the identifier of the link from which the LSA was received.
Туре	Shows the type of the link-state advertisement. Each link state type has a separate advertisement format.

Name	Description
	Link-scope LSAs not recognized by the router may be stored in the database.
RouterId	Shows the 32 bit number that uniquely identifies the originating router in the autonomous system.
Lsid	Identifies the piece of the routing domain that is being described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Indicates the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

Adding an NBMA neighbor

Add an NBMA neighbor for each interface that is eligible to become the DR.

An NMBA interface with a positive nonzero router priority is eligible to become the DR for the NBMA network and is configured with the identification of all attached routers, IPv6 addresses, and router priorities.

Before you begin

- Identify the following information:
 - specific interfaces to include in the NBMA network
 - the IPv6 address for each interface
 - the router priority for each interface
 - the hello interval for the network
 - the router dead interval for the network
 - the poll interval for the network
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.

About this task

In contrast to a broadcast network where switches multicast (send to AllSPFRouters and AllDRouters) certain OSPF protocol packets, switches replicate and send NBMA packets to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination addresses AllSPFRouters and AllDRouters. Because the NBMA network does not broadcast, you must

manually configure a list of neighbors and priorities for all routers in the network that can become the DR. Potential DRs use a positive nonzero router priority.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.
- 3. Click the **NBMA Neighbors** tab.
- 4. Click Insert.
- 5. Select the IPv6 port or VLAN interface.
- 6. Specify the IPv6 address for the neighbor.
- 7. Specify the priority for the neighbor.
- 8. Click Insert.

NBMA Neighbors field descriptions

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

Name	Description
IfIndex	Specifies the link ID for the link over which the switch reaches the neighbor.
Address	Specifies the IPv6 address of the neighbor.
Priority	Specifies the priority to use for this neighbor in the designated router election process. A value of 0 indicates the neighbor cannot become the designated router. The higher the priority value, the higher chance the switch will win the election process. The default is 1.
Rtrld	Identifies the neighboring router in the autonomous system. The value is 0.0.0.0 until the switch receives a hello message from the neighbor.
State	Identifies the state of the relationship with the neighbor. The state can be one of the following:
	• down
	attempt
	• init
	• twoWay
	exchangeStart
	exchange
	loading
	• full

Configuring Route Redistribution to OSPFv3

Configure a redistribute entry to announce routes of a certain source protocol into the OSPFv3 domain. Optionally, use a route policy to control the redistribution of routes.

About this task



Important:

Changing the OSPFv3 redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform this procedure. It is recommended that is you want to change the default preferences for an OSPFv3 redistribute context, you must do so before you enable the protocols.

Before you begin

- Enable OSPFv3 globally.
- Ensure that a route policy exists if you intend to use a route policy.
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance. The VRF must have an RP trigger of OSPFv3. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- Click OSPFv3.
- 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. Click Insert.

Redistribute Field Descriptions

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

Name	Description
DstVrfld	Specifies the destination virtual router forwarding instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.

Name	Description
SrcVrfld	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.
	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.

Modifying an OSPFv3 interface

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface. An IPv6 interface can be a tunnel, port or VLAN.

Before you begin

• The OSPFv3 interface must exist.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPFv3.
- 3. Click the Interfaces tab.
- 4. Double-click a cell to edit the value.
- 5. Click Apply.

Interfaces field descriptions

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

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
Areald	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Туре	Specifies the OSPFv3 interface type as one of the following:
	• broadcast
	• NBMA
	• point-to-point
	• point-to-multipoint
	• passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an interal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election.
	A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router.
	The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this field to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are

Name	Description
	received. You must configure this field to the same value for all routers attached to a common network. The default is 40.
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following:
	• down
	loopback
	waiting
	pointToPoint
	designatedRouter
	backupDesginatedRouter
	otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100.
	Note:
	If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.
BfdEnable	Enable Bidirectional Forwarding Detection (BFD) for OSPF.

Viewing OSPFv3 neighbors

View OSPFv3 neighbors to see routers with interfaces to a common network.

The OSPFv3 hello protocol maintains and dynamically discovers neighbor relationships.

The exception is an NBMA network; you manually configure permanent neighbors on each router eligible to become the designated router (DR).

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPFv3.
- 3. Click the **Neighbors** tab.

Neighbors field descriptions

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

Name	Description
IfIndex	Displays the local-link ID of the link over which the neighbor can be reached.
Rtrld	Identifies the neighboring router in the Autonomous System.
	The value is the router ID of the neighboring router, which in OSPF uses the same format as an IPv6 address but identifies the router independent of IPv6 address.
Address	Displays the IPv6 address for the neighbor associated with the local link.
Options	Displays the bit mask that corresponds to the options field on the neighbor.
State	Displays the state of the relationship with the neighbor.
	The value can be one of the following:
	• down
	• attempt
	• init
	• twoWay
	exchangeStart
	• exchange
	• loading
	• full
Nbrlfld	Displays the interface ID that the neighbor advertises in its hello packets on this link.
DeadIntCnt	Displays the Dead interval Count or TTL (time to live) field that indicates how many seconds remain before the system declares the Neighbor down.
	The starting value is the Router Dead Interval value and it decrements to 0 if no Hello is received for that

Name	Description
	neighbor within the interval. If no Hello is received within the interval, then the system declares the neighbor down.
	When a hello is received for the neighbor, the system resets the value to the Router Dead Interval value.

Viewing virtual neighbors

View information about the neighbors on the virtual link to the OSPFv3 backbone.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPFv3.
- 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	Shows the ID for the transit area.
Rtrld	Shows the ID for the neighboring router in the autonomous system.
Localifindex	Shows the local interface ID for the virtual link over which the switch can reach the neighbor.
AddressType	Shows the type of address as one of the following:
	• ipv4
	• ipv6
	• ipv4z
	• ipv6z
	• dns
	ipv4z and ipv6z indicate a scope zone.
Address	Shows the IPv6 address that this virtual neighbor advertises. This value must be a global scope address.
Options	Shows a bit mask that corresponds to the OSPF options field of the neighbor.

Table continues...

Name	Description
State	Shows the state of the virtual neighbor relationship. The value can be one of the following:
	• down
	attempt
	• init
	• twoWay
	exchangeStart
	exchange
	• loading
	• full

Chapter 8: RIPng

Table 19: RIPng product support

Feature	Product	Release introduced	
For configuration details, see Config	For configuration details, see Configuring IPv6 Routing for VOSS.		
RIPng	VSP 4450 Series	VOSS 5.0	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 5.0	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 5.0	
	VSP 8400 Series	VOSS 5.0	
	VSP 8600 Series	VSP 8600 6.2	
	XA1400 Series	Not Supported	

RIPng fundamentals

Routing Information Protocol next generation (RIPng) allows routers to exchange information for computing routes through an IPv6–based network. You should implement RIPng only on routers. IPv6 provides neighbor router information required by RIPng protocol to function as intended. A RIPng router is assumed to have interfaces in several networks and the protocol relies primarily on the metric of each network to compute routes using the distance vector algorithm.

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.

RIPng-enabled routers use UDP port 521 (the RIPng port) to exchange routing information. RIPng responds to a request by sending a message to the port from which the request originates. Specific queries can be sent from ports other than the RIPng port, but they must be directed to the RIPng port on the target machine.

Each router advertises routing information by sending an update every 30 seconds (one interval). If RIPng 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.

Note:

These time interval values are default values which are configurable by the user.

Each router that implements RIPng contains a routing table. This table contains one entry for every destination that is reachable throughout the system operating RIPng. At a minimum, each routing table entry contains the following information:

- The IPv6 prefix of the destination.
- A metric that represents the total cost of getting a datagram from the router to that destination. The metric is the sum of the costs of traversing the networks to arrive at the destination.
- The IPv6 address of the next router in the path to the destination (the next hop). The next-hop IPv6 address is a linklocal address.
- The VLAN or brouter port on which the RIPng routes were learned.
- The age of the RIPng route.

RIPng protocol implementation is specified in IETF document RFC 2080.

RIPng messages and packet format

RIPng-enabled routers use UDP port 521 (the RIPng port) to send and receive datagrams.

The following figure shows the RIPng packet format:

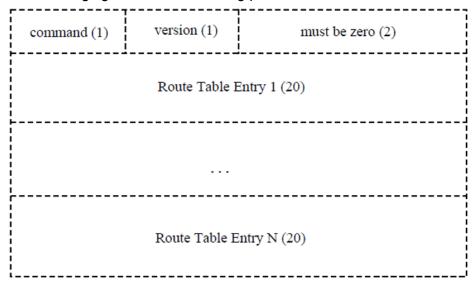


Figure 10: RIPng packet format

A RIPng packet header consists of the following components:

- **command**: Specifies the purpose of the message.
- version: The version of RIPng.

Originate Default route

Generally you use a default route when it is not convenient to list every possible network in RIPng updates, and one or more routers in the system are able to handle traffic to networks that RIPng does not explicitly list.

RIPng is enabled with the default route only option. When you enable default route only on an interface, it suppresses all other routes in the update sent for the interface, and advertises only the default route.

Timers

RIPng states four different timer intervals for protocol operation:

- Update timer: The RIPng process sends a complete routing table to each neighboring router every 30 seconds. To prevent collisions on broadcast networks, the process adds an offset value to the timer.
- **Timeout time interval**: This is a 180 second time interval associated with every route. If the time interval expires, the metric for this route updates to the value of infinity (16) and the route is no longer valid. However, the routing table retains the value for another 120 seconds.
- Garbage collection time interval: After the timeout time interval expires and the route becomes invalid, it remains in the routing table until the garbage collection time interval expires. The garbage collection time interval is 120 seconds. Until the garbage collection time interval expires all updates sent by this router include the invalid route. When the garbage collection timer expires, the process removes the route from the routing table.
- **Triggered update time interval**: The triggered update time interval is set to a random value between 1 and 5 seconds after a triggered update is sent. A single update is sent even if multiple triggered updates occur before the timer expires.

Configuration of timers or time intervals is supported only at the CLI/SNMP/EDM level. Configuration of timers or time intervals is not supported at the interface/port level.

RIPng Configuration using CLI

Use the procedures in this section to configure RIPng using CLI.

Configuring RIPng globally

Configure RIPng parameters on the router so you can control RIPng behavior on the system.

Before you begin

You can configure RIPng only on a global router. You cannot configure RIPng on a VRF instance.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RIPng globally:

```
router rip ipv6-enable
```

Configuring RIPng on an interface

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

About this task

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

Before you begin

- · Assign an IP address to the port or VLAN.
- Configure RIPng and enable it globally.

Procedure

1. Enter Interface Configuration mode:

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



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

2. Create a RIPng interface:

```
ipv6 rip
```

3. Enable the RIPng interface:

```
ipv6 rip enable
```

4. Verify the operational status of the RIPng interface:

```
show ipv6 rip interface
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #interface vlan 22
Switch:1(config-if) #ipv6 rip
Switch:1(config-if) #ipv6 rip enable
Switch:1(config-if) #show ipv6 rip interface

Total RIPng interfaces: 2

RIPng Interface - GlobalRouter

RIPng Interface - GlobalRouter

TFINDX COST POISON SEND ADMIN OPER
STATUS DEFAULT STATUS STATUS

257 (2/2 ) 2 disable disable enable enable
```

```
2070 (22 ) 5 disable disable enable disable
2 out of 2 Total Num of RIPng interfaces displayed
```

Variable definitions

Use the data in the following table to use the ipv6 rip 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. This variable applies only to VLAN interfaces, not ports.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring RIPng custom values

Configure custom values for RIPng parameters to replace default values.

Before you begin

· Configure RIPng and enable it globally.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Enable RIPng poison:

```
ipv6 rip poison enable
```

3. Specify the RIPng cost:

```
ipv6 rip cost <1-15 Cost>
```

4. Access router RIPng configuration mode:

```
router rip
```

5. Specify the RIPng holddown timer value:

```
ipv6 timers basic holddown <0-360>
```

6. Specify the RIPng timeout timer value:

```
ipv6 timers basic timeout <15-259200>
```

7. Specify the RIPng update timer value:

```
ipv6 timers basic update <1-360>
```

8. Specify the default route metric value:

```
ipv6 default-information metric <1-15)</pre>
```

9. Enable default information globally:

```
ipv6 default-information enable
```

10. Ensure the configuration is correct:

```
show ipv6 rip
```

Example

Configure custom values for RIPng.

```
Switch: 1>enable
Switch: 1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #router rip
Switch:1(config-rip) #ipv6 default-information metric 1
Switch:1(config-rip) #ipv6 default-information enable
Switch:1(config-rip) #ipv6 timers basic update 30
Switch:1(config-rip) #ipv6 timers basic timeout 180
Switch:1(config-rip) #ipv6 timers basic holddown 120
Switch:1(config-rip) #router rip ipv6-enable
Switch:1(config) #show ipv6 rip
                        RIPng Global - GlobalRouter
                 Rip : Enabled
       HoldDown Time : 120
     Timeout Interval: 180
          Update Time: 30
  Default Info Metric : 1
   Default Info State : Enabled
Default Import Metric : 1
```

Variable definitions

Use the data in the following table to use the ipv6 rip poison, ipv6 default-information and ipv6 timers basic commands.

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.
	This variable applies only to VLAN interfaces, not ports.
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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
poison enable	Enables Poison Reverse. If you disable Poison Reverse (no poison enable). 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.
<1-15 Cost>	Configures the RIPng cost for this port (link).
holddown <0-360>	Configures the RIPng holddown timer value, the length of time (in seconds) that RIPng continues to advertise a network after it determines that the network is unreachable. The default is 120.
timeout <15-259200>	Configures the RIPng timeout interval. The default is 180.
update <1–360>	Configure the RIPng update timer. The update time is the time interval between RIPng updates.
default-information <1–15>	Configure the default route metric value.

Configuring RIPng route distribution

Configure a redistribute entry to announce certain routes into the RIPng domain, including static routes, direct routes, Open Shortest Path First (OSPFv3), IS-IS, or Border Gateway Protocol (BGP +).

Before you begin

• Enable RIPng globally.

Procedure

1. Enter Global Configuration mode:

enable

```
configure terminal
```

2. Access router RIP configuration mode:

```
router rip
```

3. Enable the redistribution:

```
ipv6 redistribute {direct|isis|static|ospf|bgp} enable
```

4. Ensure the configuration is correct:

```
show ipv6 rip redistribute
```

Example

Enable the redistribution instance.

```
Switch: 1#enable
Switch: 1#configure terminal
Switch:1(config) #router rip
Switch:1(config-rip) #ipv6 redistribute bgp enable
Switch:1(config-rip)#ipv6 redistribute direct enable
Switch:1(config-rip) #ipv6 redistribute isis enable
Switch:1(config-rip) #ipv6 redistribute ospf enable
Switch:1(config-rip)#ipv6 redistribute static enable
Switch:1(config-rip) #show ipv6 rip redistribute
______
                    RIPng Redistribute List
______
     direct
                             : enabled
                              : enabled
      static
      ospf
                              : enabled
      bgp
                              : enabled
     isis
                             : enabled
```

Variable definitions

Use the data in the following table to use the ipv6 redistribute command.

Variable	Value
 	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, ospf, or static.

RIPng Configuration using EDM

Use the procedures in this section to configure RIPng using EDM.

Configuring RIPng globally

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

About this task

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

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

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click IPv6 RIPng.
- Click the Globals tab.
- 4. Select the **enable** option button.
- 5. Configure other global RIPng parameters as required.
- 6. Click Apply.

Globals field descriptions

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

Name	Description
AdminState	Enables or disables RIPng globally. The default is disabled.
UpdateTime	Specifies the time interval between RIPng updates for all interfaces. The default is 30 seconds, and the range is 1–360.
GlobalHoldDownTime	Configures the length of time that RIPng continues to advertise a network after the network is unreachable. The range is 0–360 seconds. The default is 120 seconds.
GlobalTimeOutInterval	Configures the RIPng timeout interval. The range is 15–259200 seconds. The default is 180 seconds.
DefaultInfoMetric	RIPng default-information metric.
DefaultInfoState	Default-information enable or disable at the global level
DefaultImportMetric	Specifies the RIPng default import metric.

Configuring an IPv6 RIPng interface

Configure RIPng parameters on an interface so you can control RIPng behavior on the interface.

About this task

RIPng does not operate on an interface until you enable it globally and on the interface.

You can also configure an IPv6 RIPng interface for a brouter port by selecting **Device Physical View**, selecting a port, and following the **Edit** > **Port** > **IPv6** navigation path. You can configure an IPv6 RIPng interface for a VLAN through the **VLAN** > **VLANs** > **Basic** > **IPv6** navigation path. This procedure uses the main IPv6 RIPng navigation path where you can create both types of interfaces.

Before you begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIPng globally.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6 RIPng.
- 3. Click the Interfaces tab.
- 4. Click Insert.
- 5. In the **IfIndex** box, type a value to identify the IPv6 interface.
- 6. In the RipAdminStatus option box, select enable.
- 7. Configure other parameters as required.
- 8. Click Insert.

Interfaces field descriptions

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

Name	Description
IfIndex	RIPng interface index.
RipAdminStatus	Enable or disable RIPng on an interface.
DefaultInfoState	Enable or disable default information at the interface level.
Cost	Specifies the RIPng metric cost.
Poison	Enable or disable poison reverse on an RIPng interface.
RipOperStatus	Enable or disable the RIPng operational state on an interface.

Configuring an IPv6 RIPng VLAN interface

Configure RIPng parameters on a VLAN interface so you can control RIPng behavior on the interface.

About this task

RIPng does not operate on an interface until you enable it globally and on the interface.

You can also configure an IPv6 RIPng interface for a brouter port by selecting **Device Physical View**, selecting a port, and following the **Edit** > **Port** > **IPv6** navigation path.

Before you begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.

• Enable RIPng globally.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration >VLAN**.
- 2. Click VLANs.
- 3. Click the **Basic** tab.
- 4. Select a row, and click IPv6.
- 5. Click Insert.
- 6. Configure other parameters, as required.
- 7. Click Insert.
- 8. Click Apply.

IPv6 Interfaces VLAN field descriptions

Use the data in the following table to use the **IPv6 Interfaces** tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	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 value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Туре	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.

Table continues...

Name	Description
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN.
	You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
	Different hardware platforms support different MAC offset ranges.
ForwardingEnabled	Indicates whether IPv6 forwarding is enabled.
	The default is enabled.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).

Configuring an IPv6 RIPng brouter port interface

Configure RIPng parameters on an interface so you can control RIPng behavior on the interface.

About this task

RIPng does not operate on an interface until you enable it globally and on the interface.

Before you begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- · Assign an IP address to the interface.
- Enable RIPng globally and on the interface.

Procedure

- 1. In the Device Physical View, select a port.
- 2. In the navigation pane, expand the **Configuration** >**EditPort** folders.
- 3. Click IPv6.
- 4. Click the **RIPng** tab.
- 5. In the RipAdminStatus option box, select enable.
- 6. Configure other parameters as required.
- 7. Click Apply.

RIPng field descriptions

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

Name	Description
RipAdminStatus	Enable or disable RIPng on an interface.
DefaultInfoState	Enable or disable default information at the interface level. The default is disable.
Cost	Specifies the RIPng metric cost. The default is 1.
Poison	Enable or disable poison reverse on an RIPng interface. The default is disable.
RipOperStatus	Shows the RIPng operational state on an interface.

Graphing IPv6 RIPng statistics

Use the following procedure to graph RIPng statistics for monitoring RIPng behavior on the interface.

About this task

RIPng does not operate on an interface until you enable it globally and on the interface.

Before you begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIPng globally and on the interface.

Procedure

- 1. In the navigation pane, expand the **Configuration >lpv6** folders.
- 2. Click the IPv6 RIPng.
- 3. Click the Stats tab.
- 4. Select an interface row.

- 5. Click Graph.
- 6. Click **Apply**.

Configuring route redistribution to RIPng

Configure a redistribute entry to announce routes of a certain source protocol type into the RIPng 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.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click IPv6 RIPng.
- 3. Click the **Redistribute** tab.
- 4. Double-click the value in the **Enable** column that corresponds with the source protocol type you want to enable or disable.
- 5. Select **enable** or **disable** from the list.
- 6. Click Apply.

Redistribute field descriptions

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

Name	Description
DstVrfild	Specifies the destination VRF ID used in the redistribution.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrfld	Specifies the source VRF ID used in the redistribution.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) a RIPng redistribute entry for a specified source type.

Viewing stats for RIPng interfaces

View statistics for RIPng interfaces.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IPv6**.

- 2. Click IPv6 RIPng.
- 3. Click the **Stats** tab.

Stats field descriptions

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

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface.
RcvBadPackets	The number of RIPng response packets received by the RIP process that were subsequently discarded for any reason (examples: a version 0 packet or an unknown command type).
RcvBadRoutes	The number of routes, in valid RIPng packets, that were ignored for any reason (examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIPng updates actually sent on this interface.
RcvUpdates	The number of triggered RIPng updates actually received on this interface. This explicitly does not include full updates received containing new information.

Chapter 9: VRRP

Table 20: Virtual Router Redundancy Protocol for IPv6 product support

Feature	Product	Release introduced
For configuration details, see	<u>S</u> .	
IPv6 VRRP	VSP 4450 Series	VOSS 4.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 4.1
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported
VRRPv3 for IPv6	VSP 4450 Series	VOSS 5.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 5.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 5.1
	VSP 8400 Series	VOSS 5.1
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported

VRRP

For IPv6 hosts on a LAN to learn about one or more default routers, IPv6-enabled routers send router advertisements using the IPv6 ND protocol. The routers multicast these router advertisements every few minutes.

The ND protocol uses a mechanism called neighbor unreachability detection to detect the failure of a neighbor node (router or host) or the failure of the forwarding path to a neighbor. Nodes can monitor the health of a forwarding path by sending unicast ND neighbor solicitation messages to the neighbor node. To reduce traffic, nodes only send neighbor solicitations to neighbors to which they

actively send traffic and only after the node receives no positive indication that the neighbors are up for a period of time. A host takes a minimum of 5 seconds to learn that a router is unreachable before it switches to another default router, but this minimum value increases ND traffic. This delay can cause service disruption.

VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol. With VRRP for IPv6, a backup router can take over for a failed default router in approximately three seconds (using default parameters). The switchover is accomplished without interaction with the hosts and with a minimum amount of VRRP traffic.

The IPv6 VRRP implementation is similar to the existing IPv4 VRRP operation, including support for holddown timer, critical IP, fast advertisements, and backup master. With backup master enabled, the backup switch routes all traffic according to its routing table. The backup master switch does not perform Layer 2 switching for the traffic to the VRRP master.

You must specify a link-local address to associate with the virtual router. Optionally, you can also assign global unicast IPv6 addresses to associate with the virtual router. Network prefixes for the virtual router are derived from the global IPv6 addresses assigned to the virtual router.

One active master switch exists for each IPv6 network prefix. All other VRRP interfaces in a network are in backup mode.

VRRP for IPv6 operation

VRRP uses a virtual IP address shared between two or more routers connecting the common network prefix to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides dynamic default gateway redundancy in the event of failover.

The VRRP router with higher priority is called the master router. In case of equal priority the router with higher link-local address becomes the master router. The master router forwards packets sent to the virtual router IP addresses.

The following figure shows the minimum VRRP topology.

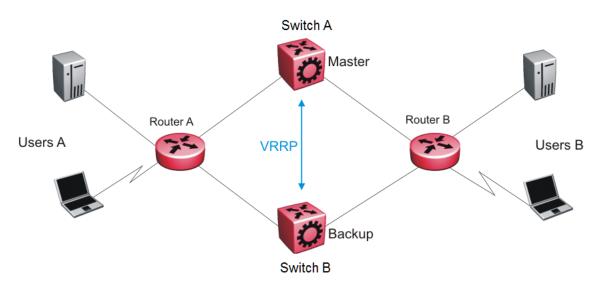


Figure 11: VRRP network topology

Traffic flows between users A and users B.

Router A uses VRRP global addresses as next hops for users B, and Router B for users A.

The VRRP master forwards the traffic and sends VRRP advertisements in the VLAN to announce to the backups that it is the master. If the master is no longer available, the backup takes over and becomes master. The only change occurs to the state of VRRP.

The VRRP router then transitions to the controlling state.



The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. The router responds to ND neighbor solicitation and ND router solicitation messages for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts packets addressed to IP addresses associated with the virtual router.

If you initialize the VRRP router and the priority is not 255, the router transitions to the backup state to ensure that all Layer 2 switches in the downstream path relearn the new origin of the VRRP MAC addresses.

In the backup state, a VRRP router monitors the availability and state of the master router. The backup does not respond to ND neighbor solicitation and ND router solicitation messages for virtual router IP addresses and discards packets with a MAC address equal to the virtual router MAC address. The backup does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, it transitions back to the initialize state. If the master router goes down, the backup router sends the VRRP advertisement and unsolicited ND neighbor advertisements and ND router advertisements described in the preceding paragraphs and transitions to the controlling state.

VRRP advertisements and master router failover

When you initialize a VRRP router, the master router continues to send advertisement messages at the advertisement interval period.



Note:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

The other VRRP routers transition to the backup state in the following situations:

- if the priority in the received advertisement is greater than the local priority
- if the priority in the received advertisement is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address

The backup routers use the advertisements from the master router as a keepalive to monitor the health of the master router. If the backup router does not receive an advertisement during the master downtime interval, calculated as 3 * advertisement interval, then the master router is declared down.

If a shutdown occurs, the master router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state

The priority value 0 indicates that the master router has stopped participating in VRRP. This value triggers the backup router to transition to the master state without waiting for the current master to time out.

Critical IPv6 address and holddown timer

The critical IPv6 address is an interface that has primary impact on VRRP. If you enable critical IPv6 and the status of the critical IP changes, the master and backup relationship also changes.

If you configure and enable critical IPv6 address, the master transitions to backup if the critical IPv6 is down, and the backup becomes the master. After the critical IPv6 address of the original master resumes, if the hold-down timer is configured to 0, it becomes the master immediately. Otherwise, the original master transitions to the master state after the hold-down timer time out.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that, if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

The critical address can be one of the global unicast IPv6 addresses assigned to any local IPv6 interfaces.

The holddown timer is a proprietary enhancement to VRRP.

After a master transitions to backup by critical IP changing, one of the backup routers will be elected as the master router. After the critical IPv6 of the original master is restored, the original master remains in the backup state for a period of time that you configure by using the holddown-timer parameter. The router becomes the master immediately if you use the command ipv6 vrrp <1-255> action preempt.

The holddown timer allows the master router enough time to detect and update the dynamic routes. The timer delays the preemption of the master over the backup, when the master becomes available. If the hold-timer is configured to 0, it becomes the master router immediately. Otherwise, it transitions to the master state only after the holddown timer times out.

The holddown timer does not apply during failovers caused by VRRP router priority change. The holddown timer applies only to failovers caused by a critical IP failure.

Configure all of your routers to use identical values for the holddown timer.

! Important:

Do not use VRRP backup master and critical IP at the same time. Use one or the other. The critical IP address must be a local address.

VRRP backup master with triangular SMLT

The standard implementation of VRRP supports one active master switch for each IPv6 subnet. All other VRRP interfaces in a network are in backup mode.

A deficiency occurs when VRRP-enabled switches use SMLT. If VRRP switches are aggregated into two SMLT switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the Multilink Trunk [MLT] traffic distribution algorithm).

However, VRRP usually has only one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, all traffic that reaches the backup VRRP router is forwarded over Virtual Inter-Switch Trunk (vIST) toward the master VRRP router. In this case, vIST potentially does not have enough bandwidth to carry all the aggregated traffic.

To resolve this issue, assign the backup router as the backup master router. The backup master router can actively load-share the routing traffic with a master router.

Because the two VRRP peer nodes exchange MAC address tables, the VRRP backup master can forward traffic directly, on behalf of the master router. The switch in the backup master state routes all traffic received on the backup master IP interface according to its routing table. The backup master switch does not perform Laver 2 switching for the traffic to the VRRP master.

If you enable SMLT on the backup master router, the incoming host traffic is forwarded over the SMLT links as usual.

Important:

Do not use VRRP backup master and critical IP at the same time. Use one or the other.

Fast advertisement

You can configure the advertisement time interval (in seconds) between sending advertisement messages. This interval permits fast network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence, which means network problems must be detected within hundreds of milliseconds.

To meet these requirements, the fast advertisement interval are provided.

The fast advertisement interval is similar to the advertisement interval parameter except for the unit of measure and the range. The fast advertisement interval is expressed in milliseconds and the range is from 200 to 1,000 milliseconds. This unit of measure must be in multiples of 200 milliseconds.

To configure fast advertisement, you must specify a fast advertisement interval and explicitly enable the fast advertisement option. After you enable fast advertisement, the fast advertisement interval is used instead of the advertisement interval.

If you enable fast advertisement, VRRP can only communicate with other products that have the same configuration.

Accept-mode

When you configure VRRP for IPv6 on an interface you can configure the accept-mode parameter. which controls whether the VRRP master or backup master accepts packets destined for the IPv6 address associated with the virtual router.

By default, accept-mode is disabled. The accept-mode parameter does not affect the Neighbor Discovery packets. The master router forwards packets with a destination link-layer MAC address that matches the virtual MAC address, and accepts packets forwarded over the virtual interswitch trunk (vIST) toward the master router, if accept-mode is enabled. If you disable accept-mode, you cannot ping the virtual IPv6 address. If you enable accept-mode, the master router accepts packets addressed to the IPv6 address that is associated with the virtual router.

When you configure VRRP for IPv6 on an interface, you can configure the accept-mode parameter. By default, accept-mode is disabled. If you disable accept-mode, the master router does not drop neighbor solicitations or neighbor advertisements. The master router forwards packets with a destination link-layer MAC address that matches the virtual MAC address. If you disable acceptmode, you cannot ping the virtual IPv6 address.

Note:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

VRRPv3

VRRPv3 is a combined protocol for both IPv4 and IPv6. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router that controls the IPv4 or IPv6 addresses associated with a virtual router is called the Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Backups wait for a Master and take ownership when the Master is no longer detected.

The election protocol provides dynamic failover in the forwarding responsibility when the Master is unavailable. VRRP for IPv4 gains a higher-availability default path without configuring dynamic routing or router discovery protocols on every end-host. VRRP for IPv6 gains a quick switch-over to Backup routers compared to the standard IPv6 Neighbor Discovery mechanisms.

Note:

The VRRP IPv6 link-local address must be the same for all VRRP routers sharing the same link and the same virtual router ID, that is, the same VRRP instance. It is the address the VRRP advertisements are sent from. Also, the Router Advertisement packets from the VRRP interface are transmitted using this address, so, when IPv6 Stateless Address Autoconfiguration is used, this address is added to host Default Router List and is used as a gateway.

The software supports VRRPv3 for IPv4 and VRRPv3 for IPv6. VRRPv3 for IPv6 is compliant to RFC 5798. The software also supports VRRPv2 for IPv4.

VRRPv3 guidelines

The switch also supports VRRPv2 for IPv4. If you configure VRRP IPv6 on an interface, it runs independently of the IPv4 version. Configure the version of the VRRP IPv4 on the interface before you configure any other IPv4 VRRP attributes. By default, the version is not configured to a particular value. However, when sourcing older configuration files that do not have the version saved, the router configures the version to VRRPv2 by default. If you change the version, all IPv4 configuration under that interface is automatically removed, and you are prompted for a confirmation before this operation.

Perform the CLI configuration through ip vrrp or ipv6 vrrp nodes; CLI commands for IPv4 are common for version 2 and version 3.

The following list identifies the features that make both IPv4 and IPv6 VRRPv3 features compliant to RFC 5798:

 Advertisement vs Fast-advertisement — Prior to RFC 5798, the minimum advertisement interval was 1 second, with Fast-advertisement a sub-second interval could be configured. When this feature is enabled, the VRRP ADVERTISEMENT packets are sent with type 7 instead of 1. With RFC 5798 the sub-second interval is standardised, and the switch sends all packets for VRRPv3 with type 1. The use of Fast-advertisement remains the same. VRRPv2 packets send with type 7, if Fast-advertisement is enabled.

- Add Master-advertisement-interval Prior to RFC 5798 compliance, all virtual routers on the same VLAN had the same Advertisement-Interval configured. RFC 5798 states that you can use different Advertisement Intervals on the Master and Backup. On the Master, the Masteradvertisement-interval and the Advertisement-Interval have the same value. On the Backup, the Master-advertisement-interval is used to calculate the timers, and the locally configured Advertisement-Interval is ignored until the Backup transitions to Master. The Masteradvertisement-interval value is put in the advertisement packet type sent by the Master
- Transition to master as specified in RFC 5798 Prior to RFC 5798, if a Backup receives an advertisement with a lower priority (or same priority but lower IP), it immediately sends its own advertisement and transitions to Master. However, RFC 5798 states that such packets must be discarded, which means it will transition to Master after the Master Down Timer expires
- Add skew-time RFC 5798 states that skew-time is calculated depending on the priority, and Master-advertisement-interval assures that the Backup with highest priority sends the first advertisement when the Master goes down

```
Skew time is calculated using the formula: (((256 - priority) *
Master Adver Interval) / 256).
```

• Add preempt-mode — Preempt-mode is different from the ipv6 vrrp <vrid> action preempt command, which is an operational command issued when you want to stop the holddown timer. RFC 5798 states that preempt-mode should be set to false when you do not want a higher priority Backup to transition to Master. By default, it is set to true



Note:

Accept-mode is not fully implemented for IPv4 VRRPv3. You can only ping the virtual IP address, the same way as it is for IPv4 VRRPv2.

VRRP Configuration using CLI

Use the procedures in this section to configure VRRP using CLI.

Configuring the VRRP interface

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts, in order to create a VRRP instance.

Before you begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.

• You must specify a link-local address to associate with the virtual router.

About this task

VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network.

VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Perform this procedure to also configure the additional addresses for which the virtual router acts as a backup.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

- 2. Associate an address with the virtual router for either link-local or global:
 - ipv6 vrrp address <1-255> link-local WORD <0-127>
 - ipv6 vrrp address <1-255> global WORD <0-255>

Note:

You must configure the link-local address before you configure the global address.

3. Enable VRRP for the interface:

```
ipv6 vrrp \langle 1-255 \rangle enable
```

Example

Associate a link-local address with the virtual router ID 12:

```
Switch:1(config-if) #ipv6 vrrp address 12 link-local fe80::1234
```

Associate a global address with the virtual router ID 12

```
Swith:1(config-if) #ipv6 vrrp address 12 global 3333::1234/64
```

Enable VRRP for the interface:

```
Switch:1(config-if) #ipv6 vrrp 12 enable
```

Variable definitions

Use the data in the following table to use the ipv6 vrrp address command.

Variable	Value
<1-255>	Specifies the virtual router ID. The virtual router acts as the default router for one or more associated addresses.
enable	Enables IPv6 VRRP. The default is disabled.
global WORD <0-255>	Specifies a global IPv6 address and mask to associate with the virtual router.
link-local WORD <0-127>	Specifies a link-local IPv6 address to associate with the virtual router.

Use the data in the following table to use the interface 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.
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing VRRP information

Display VRRP port or VLAN information to verify your configuration. Show VRRP information by IPv6 address or virtual router ID. If you enter a virtual router ID or an IPv6 address when you view VRRP information, the information applies only to that virtual router ID or for that interface.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the configuration information for all interfaces:

show ipv6 vrrp interface [verbose] [vrf WORD < 1-16 > | vrfids WORD < 0-512 >]

3. View the configuration information for one or more ports:

show ipv6 vrrp interface gigabitethernet $\{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]\}$ [verbose] [vrf WORD<1-16> | vrfids WORD<0-512>]

4. View the configuration information for one or more VLANs:

show ipv6 vrrp interface vlan [<1-4059>] [verbose] [vrf WORD<1-16> | vrfids WORD<0-512>]

5. View the configuration information for one or more virtual router IDs:

show ipv6 vrrp interface vrid <1-255> [verbose] [vrf WORD<1-16> | vrfids WORD<0-512>]

6. View VRRP address information:

show ipv6 vrrp address [vrf WORD<1-16> | vrfids WORD<0-512>]

7. View VRRP address information for a link-local address:

show ipv6 vrrp address link-local WORD < 0-127 > [verbose] [vrf WORD < 1-16 > |verbose] [vrf WORD < 0-512 >]

8. View VRRP address information for a virtual router ID:

show ipv6 vrrp address vrid <1-255> [vrf WORD<1-16> | vrfids WORD<0-512>]

Example

Switch:1>show ipv6 vrrp address										
VRRP Info - GlobalRouter										
VRID	P/V	IP			MAC			STA	ΓE	CONTROL
12	1/1	fe80:0:0	:0:0:0:0:1234		00:00:5	e:(t	
VRID		MASTER			PRIO A	DV	UP TIME	3		
12	1/1		0:0:0:0		100 1		0 day(s	s), 0	0:00:0	0
VRID	,	CRITICAL				Ι	CRITICAL ENABLED		MODE	
12		0:0:0:0:	0:0:0:0				 No			
VRID		MASTER	BACKUP-MASTER STATE	ADV	, ,			DWN		
12			down							
VRID	P/V	GLOBAL A	DDRESS							
12	1/1	1111::22	 22/64							
Flags Legend: HLD DWN: Configured hold-down timer value, REM: REMaining hold-down timer value										

More (q = quit)							
Switch:1#show ipv6 vrrp interface verbose							
Vlan Vrrp for IPv6 Extended							
VLAN VRF MASTER	VRRP						
ID NAME IPADDR	ID	STATE	CONTROL	PRIORITY			
40 GlobalRouter 40 GlobalRouter			disable disable		0:0:0:0:0:0		
All 2 Vlan Vrrp Exten	ded E	ntries o	at of 18	Total Num	of Vrrp dis	played	
VLAN VRF CRITICAL	VRRP	HOLDDWN	ACTION	CRITICAL			
ID NAME IPADDR	ID		TIME	IP ENABLE			
40 GlobalRouter 40 GlobalRouter					0:		
All 2 Vlan Vrrp Exten	ded E	ntries o	at of 18	Total Num	of Vrrp dis	played	
VLAN VRF PREEMPT	VRRP	BACKUP	BACKUP	ADVERTISE	FAST ADV	FAST ADV	MASTER ADV
ID NAME MODE	ID	MASTER	MASTER	INTERVAL	INTERVAL	ENABLE	INTERVAL
MODE			STATE	(s)	(ms)		(ms)
 40 GlobalRouter	1	disable	down	1	200	disable	1000
enable 40 GlobalRouter					200	disable	
enable	_	arbabic	aown	_	200	arbabre	1000
All 2 Vlan Vrrp Extend	ded E	ntries o	ut of 18	Total Num	of Vrrp dis	played	
=======================================						=======	
]	Port Vrrp	o for IP	76 Extended	d =======		
PORT VRF	VRRI	P					
MASTER NUM NAME IPADDR	ID	STATE	CONTRO	L PRIORITY			
1/23 GlobalRouter	1	init	disable	= 100	0:0:0:0:0:	0:0:0	
1/23 GlobalRouter 1/23 GlobalRouter	2	init	disable	e 100	0:0:0:0:0:	0:0:0	
All 2 Port Vrrp Extend	ded E	ntries o	ut of 18	Total Num	of Vrrp dis	played	
PORT VRF CRITICAL	VRRI	P HOLDDWI	N ACTION	CRITICAL			
NUM NAME IPADDR	ID		TIME	IP ENABLI	E 		
					_		

1/23 1/23	GlobalRouter GlobalRouter	1 2	0	none none	disable disable	0:0:0:0:0: 0:0:0:0:0:		
All 2	Port Vrrp Extend	led En	tries ou	t of 18	Total Num o	f Vrrp disp	layed	
PORT PREEM	VRF PT	VRRP	BACKUP	BACKUP	ADVERTISE	FAST ADV	FAST ADV	MASTER ADV
NUM MODE	NAME	ID	MASTER	MASTER	INTERVAL	INTERVAL	ENABLE	INTERVAL
(ms)				STATE	(s)	(ms)		
1/23 enabl		1	disable	down	1	200	disable	1000
	GlobalRouter	2	disable	down	1	200	disable	1000
All 2 Port Vrrp Extended Entries out of 18 Total Num of Vrrp displayed								

Variable definitions

Use the data in the following table to use the show ipv6 vrrp commands.

Variable	Value
{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 the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
link-local WORD<0-127>	Displays information by link-local IPv6 address.
verbose	Displays extended information.
vlan [<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.
vrid <1–255>	Displays information by virtual router ID.
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

Configuring VRRP notification control

Perform this procedure to configure VRRP notification control.

Before you begin

• Assign an IPv6 address to the interface.

· Enable routing globally.

About this task

By default, generation of SNMP traps for VRRP events is enabled.

Procedure

1. Enter VRRP Router Configuration mode:

```
enable
configure terminal
router vrrp
```

2. Enable the VRRP-router to generate SNMP traps for events:

```
ipv6 send-trap enable
```

Example

Disable generation of SNMP traps for VRRP events:

```
Switch:1(config-vrrp) #no ipv6 send-trap enable
```

Configuring additional VRRP parameters for an interface

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts. VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network. VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Configure the parameters in this procedure if the default values do not meet your requirements.

Before you begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.

About this task

A switch that acts as a VRRP master does not reply to SNMP get requests to the VRRP virtual interface address. The switch will, however, respond to SNMP get requests to the physical IP address.

Procedure

1. Enter Interface Configuration mode:

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

Note:

If the 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 accept mode of the master router:

```
ipv6 vrrp <1-255> accept-mode enable
```

3. Determine if the router overrides the holddown timer:

```
ipv6 vrrp <1-255> action <none|preempt>
```

4. Configure the interval between advertisement messages:

```
ipv6 vrrp <1-255> adver-int <1-40>
```

5. Enable the backup VRRP switch for traffic forwarding:

```
ipv6 vrrp <1-255> backup-master enable
```

6. Configure the IP interface on the local router:

```
ipv6 vrrp <1-255> critical-ipv6-addr WORD<0-46> [critical-ipv6 enable]
```

7. Configure the fast advertisement interval:

ipv6 vrrp
$$\langle 1-255 \rangle$$
 fast-adv enable [fast-adv-int $\langle 200-1000 \rangle$]

8. Configure the holddown timer:

9. Configure the priority for the VRRP router:

Example

Configure the fast advertisement interval:

Switch:1(config-if) #ipv6 vrrp 12 fast-adv enable fast-adv-int 400

Configure the holddown timer:

Switch:1(config-if) #ipv6 vrrp 12 holddown-timer 30

Variable definitions

Use the data in the following table to use the **ipv6 vrrp** command.

Variable	Value
<1-255>	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses.

Table continues...

Variable	Value
accept-mode enable	Controls whether the VRRP master or backup master accepts packets (other than neighbor discovery packets) destined to the IPv6 address associated with the virtual router. The default value is disable.
action <none preempt></none preempt>	Lists options to override the holddown timer manually and force preemption:
	none does not override the timer.
	preempt preempts the timer.
	This parameter applies only if the holddown timer is active.
adver-int <1-40>	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second. Only the master router sends advertisements.
backup-master enable	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the vIST. The default is disabled.
critical-ipv6 enable	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
critical-ipv6-addr WORD<0-46>	Specifies an IP interface on the local router configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.
fast-adv enable	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.
fast-adv-int <200-1000>	Configures the interval between VRRP advertisement messages. You must configure the same value on all participating routers.
	This unit of measure must be in multiples of 200 milliseconds. The default is 200.
holddown-timer <0-21600>	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.
priority <1-255>	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.

Enabling IPv6 VRRP preempt-mode

You can configure IPv6 VRRP to preempt the existing router. If a new VRRP router is added to the network with a higher priority than the existing routers, then the new router becomes the master. If preempt-mode is disabled, then the new router does not become a master, it transitions to master only when the current master is down. By default, preempt-mode is enabled.

Procedure

1. Enter Interface Configuration mode:

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



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

2. Enter the following command:

```
ipv6 vrrp <vrid> preempt-mode enable
```

3. Use the following command to set the IPv6 VRRP preempt-mode to its default value:

```
default ipv6 vrrp <vrid> preempt-mode
```

4. Use the following command to disable the IPv6 VRRP preempt-mode:

```
no ipv6 <vrid> preempt-mode enable
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Enabling IPv6 VRRP preempt-mode for interface 1/2

```
Switch:1(config-if)# ipv6 vrrp 1 preempt-mode enable
```

Variable definitions

Use the data in the following table to use the ipv6 vrrp <vrid> command.

Variable	Value
enable	Enables preempt-mode for VRRPv3 for IPv6.
vrid <1–255>	Specifies the virtual router ID.

VRRP Configuration using EDM

Use the procedures in this section to configure VRRP using EDM.

Configure VRRP for an Interface

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts. VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network. VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Perform this procedure to configure VRRP on either a brouter port or a VLAN.

Before you begin

- · Assign an IPv6 address to the interface.
- · Enable routing globally.
- · Do not configure RSMLT on the VLAN.
- Change the VRF instance as required to configure a VRRP for an interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
- 2. Select VRRP.
- 3. Select the following tab:
 - For VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, and VSP 8400 Series, select V3 Interface.
 - For VSP 8600 Series, select Interface.
- 4. Select Insert.
- 5. Beside the **IfIndex** field, click **Port** or **VLAN**.
- 6. In the dialog box that appears, select a port or VLAN.
- 7. Select OK.
- 8. Type the virtual router ID.
- To control the packets sent to the IPv6 address associated to the virtual router, select the AcceptMode check box.
- 10. Type the primary IP address.
- 11. Select Insert.

V3 Interface or Interface Field Descriptions

Use the data in the following table to use the V3 Interface or Interface tab.

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
InetAddrType	Specifies the address type for the VRRP interface.
Vrld	Specifies a number that uniquely identifies a virtual router on a VRRP router.
PrimarylpAddr	Specifies the link-local address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Shows the state of the virtual router interface. The possible states are
	initialize—waiting for a startup event
	backup—monitoring availability and state of the master router
	master—functioning as the forwarding router for the virtual router IP addresses
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
Advinterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second.
MasterlpAddr	Specifies the IP address of the physical interface of the Master's virtual router.
UpTime	Indicates the time interval since this virtual router exited the INIT state.
CriticallpAddr	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CriticallpAddrEnabled	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
BackUpMaster	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the vIST. The default is disabled.

Table continues...

Name	Description
BackUpMasterState	Indicates if the Backup-Master is up. If the switch is in Master state, but Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
FasterAdvintervalEnabled	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.
FasterAdvInterval	Configures the interval between VRRP advertisement messages. The default is 200. Enter the values in multiples of 200 milliseconds.
AcceptMode	Controls whether the VRRP master or backup master accepts packets (other than neighbor discovery packets) destined to the IPv6 address associated with the virtual router. The default value is disable.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master. The default is enabled.
Action	Lists options to override the hold-down timer manually and force preemption:
	none does not override the timer.
	preemptHoldDownTimer preempts the timer.
	This parameter applies only if the holddown timer is active.
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
MasterAdvinterval	On the VRRP master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

Configuring VRRP notification control

Perform this procedure to configure VRRP notification control.

Before you begin

- Assign an IPv6 address to the interface.
- Enable routing globally.

• Change the VRF instance as required to configure VRRP notification control on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: Configuration > IPv6.
- 2. Click VRRP.
- 3. Click the Globals tab.
- 4. Select enabled.
- 5. Click Apply.

Globals field descriptions

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

Name	Description
NotificationCntrl	Indicates whether the VRRP-enabled router generates SNMP traps for events.
	enabled: Generate SNMP traps.
	disabled: Do not generate SNMP traps.
	The default is enabled.

Configuring additional addresses on the VRRP brouter port

Perform this procedure to configure the additional addresses for which the virtual router acts as a back up.

Before you begin

- · Assign an IPv6 address to the interface.
- Enable routing globally.

Procedure

- In the navigation pane, expand the following folders: Configuration > Edit > Port.
- 2. Click IPv6.
- 3. Click the VRRP tab.
- 4. Click AssociatedIPAddr.
- 5. Click Insert.
- 6. Type the address.
- 7. Type the prefix length.
- 8. Click Insert.

Configuring additional addresses on the VRRP interface

Perform this procedure to configure the additional addresses for which the virtual router acts as a back up.

Before you begin

- · Assign an IPv6 address to the interface.
- · Enable routing globally.
- Do not configure RSMLT on the VLAN.
- Change the VRF instance as required to configure additional addresses on the VRRP interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click VRRP.
- 3. Click the **Interface** tab.
- 4. Select an interface.
- 5. Click Associated IPAddr.
- 6. Click Insert.
- 7. Type the address.
- 8. Type the prefix length.
- 9. Click Insert.

Address List field descriptions

Use the data in the following table to use the **Address List** tab.

Name	Description
IpAddr	Specifies an IP address that is associated with a virtual router. The number of rows on this tab equals the number of IP addresses associated (backed up) by the virtual router
IpAddrPrefixLength	Specifies the length of the prefix in bits.

Chapter 10: RSMLT

Table 21: RSMLT for IPv6 product support

Feature	Product	Release introduced	
For configuration details, see Config	guring IPv6 Routing for VOSS.		
IPv6 RSMLT	VSP 4450 Series	VOSS 4.1	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 4.1	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 6.2	
	XA1400 Series	Not Supported	

RSMLT

Routed Split Multi-Link Trunking (RSMLT) is an enhancement to SMLT that enables the exchange of Layer 3 information between peer nodes in a switch cluster. RSMLT provides two main advantages over SMLT:

- provides backup for the peer after the peer goes down
- routes traffic on behalf of the peer to prevent Virtual Inter-Switch Trunk (vIST) overload

IPv6 RSMLT enables the subsecond failover for IPv6 forwarding.

The overall model for IPv6 RSMLT is essentially identical to that of IPv4 RSMLT. In short, RSMLT peers exchange their IPv6 configuration and track their states by using vIST messages. An RSMLT node always performs IPv6 forwarding on the IPv6 packets destined to the MAC addresses of the peer. If an RSMLT node detects that the RSMLT peer is down, the node forwards IPv6 traffic destined to the IPv6 addresses of the peer.

With RSMLT enabled, an SMLT switch performs IP forwarding on behalf of the SMLT peer, which prevents IP traffic from being sent over the vIST.

IPv6 RSMLT supports the full set of topologies and features supported by IPv4 RSMLT, including SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs.

Because you configure RSMLT on a VLAN, not at the IP layer, the configuration applies to both IPv4 and IPv6. You cannot enable or disable RSMLT on a VLAN for IPv6 but not IPv4; or for IPv4 but not IPv6.

With IPv6, you must configure the RSMLT peers to use the same set of IPv6 prefixes.

Supported routing protocols include the following:

- IPv6 static routes
- OSPFv3

Note:

IPv6 RSMLT is not virtualized, therefore it is not possible to enable IPv6 and RSMLT together on a VLAN which is associated with a VRF.

Configuration errors will appear if you attempt to perform the following:

- Create an IPv6 interface on a VLAN which is associated with a VRF and RSMLT is enabled on the VLAN.
- Enable RSMLT on an IPv6 enabled VLAN which is associated with a VRF.

For more information about the IPv4 RSMLT model, see <u>Configuring IPv4 Routing for VOSS</u>. This section focuses on the differences between the IPv4 and IPv6 models.

IPv6 differences

The following list identifies ways in which the IPv6 implementation of RSMLT differs from the IPv4 implementation of RSMLT.

- After the switch begins to forward traffic on behalf of the peer, duplicate address detection (DAD) is not executed for the IPv6 address of the peer. The implementation assumes that the peer IPv6 address is already known to be unique.
- An RSMLT switch installs a neighbor entry for the peer IPv6 address immediately after the peer disappearance is detected, possibly while a route for the peer still exists. This action can result in packets destined to the peer IPv6 address being delivered to the CP for a short period of time.
- You cannot configure a vIST with IPv6 peer address
- In a dual-stack VLAN, adding or deleting IPv4 or IPv6 does not affect the RSMLT functionality
 of one another. If you add IPv4 or IPv6 to an existing IPv6 or IPv4 RSMLT VLAN, the RSMLT
 state for the protocol you add second will be the same as the previous RSMLT state.

RSMLT Configuration using CLI

Use the procedures in this section to configure RSMLT using CLI.

Configuring RSMLT on a VLAN

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster. This configuration applies to both IPv4 and IPv6.

Before you begin

- An IP routing protocol is enabled on VLAN Layer 3 interfaces.
- VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About this task

The VLAN can be either IPv4 or IPv6, or both. RSMLT configuration on a VLAN simultaneously affects both IPv4 and IPv6. By default, RSMLT is disabled on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

2. Configure the holddown timer:

```
ip rsmlt holddown-timer <0-3600>
```

3. Configure the holdup timer:

```
ip rsmlt holdup-timer <0-9999>
```

4. Enable RSMLT on the VLAN:

```
ip rsmlt
```

Example

Configure the holddown timer:

```
Switch:1(config-if) #ip rsmlt holddown-timer 100
```

Configure the holdup timer:

```
Switch:1(config-if) #ip rsmlt holdup-timer 200
```

Enable RSMLT on the VLAN:

```
Switch:1(config-if) #ip rsmlt
```

Variable definitions

Use the data in the following table to use the <code>ip rsmlt</code> command.

Variable	Value
holddown-timer <0-3600>	Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The default is 60.
	If you disable RSMLT on a VLAN, non default values for this variable do not save across restarts.
holdup-timer <0-3600 9999>	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.
	If you disable RSMLT on a VLAN, non default values for this variable do not save across restarts.

Use the data in the following table to use the interface 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.

Enabling RSMLT Edge support

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

About this task

RSMLT Edge support configuration applies to both IPv4 and IPv6. You do not configure IPv4 and IPv6 separately.

The RSMLT Edge support default is disabled.

Procedure

1. Enter Global Configuration mode:

enable
configure terminal

2. Enable RSMLT Edge support:

ip rsmlt edge-support

Example

If you have enabled RSMLT Edge Support, disable the feature as follows:

Switch:1(config) #no ip rsmlt edge-support

Viewing RSMLT information

Show RSMLT information to view data about all RSMLT interfaces. The output of the command includes IPv6 information for the local and peer nodes.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Show RSMLT information about the interface:

```
show ip rsmlt [local|peer] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

3. View the status of the switch to act as a peer forwarder:

show ip rsmlt edge-support

Example

Switch:1>show ip rsmlt						
=====		Ip Rsmlt Local In				
VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
101 102	101.1.1.32 102.1.1.32	00:24:7f:9e:da:01 00:24:7f:9e:da:02	Enable Enable	Up Up	100 60	200 180
VID	SMLT ID					
	101 102					
VID	IPv6	MAC	ADMIN		HDTMR	HUTMR
101	1010:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:				100	200
102	1020:0:0:0:0:0:0:	00:24:7f:9e:da:02	Enable	Up	60	180
	1020:0:0:0:0:0:0: fe80:0:0:0:224:	0:32/64 7fff:fe9e:da02/128				
VID	SMLT ID					
101 102	101 102					

		Ip Rsm	nlt Peer Inf	o - Glob	alRout	er	
VID	IP	MAC		ADMIN	OPER	HDTMR	HUTMR
101 102	101.1.1.33 102.1.1.33	00:24:7 00:24:7	f:9e:ea:01 f:9e:ea:00	Enable Enable	Up Up	100 60	200 180
VID	HDT REMAIN	HUT REMAIN	SMLT ID				
101 102	60 60	180 180					
VID	IPv6	MAC		ADMIN	OPER	HDTMR	HUTMR
101	1010:0:0:0: 1010:0:0:0: fe80:0:0:0:	0:0:0:0/64	f:9e:ea:01	Enable	Up	100	200
102	1020:0:0:0: 1020:0:0:0:	00:24:7	f:9e:ea:00	Enable	Up	60	180
VID	HDT REMAIN	HUT REMAIN	SMLT ID				
101 101	60	180					
102	60	180					
Switc	ch:1>show ip	rsmlt edge-s	support				
RSMLT	Peer Info: rsmlt-pee	r-forwarding	; disable				

Variable definitions

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

Variable	Value
local	Shows local RSMLT information.
peer	Shows RSMLT information for the peer.
vrf WORD<1-16>	Shows information for a specific VRF name.
vrfids WORD<0-512>	Shows information for a specific VRF ID.

RSMLT Configuration using EDM

Use the procedures in this section to configure RSMLT using EDM.

Configuring RSMLT on a VLAN

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster. This configuration applies to both IPv4 and IPv6.

Before you begin

- Enable an IP routing protocol on VLAN Layer 3 interfaces.
- Ensure VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About this task

The VLAN can be either IPv4 or IPv6, or both. RSMLT configuration on a VLAN simultaneously affects both IPv4 and IPv6.

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.
- Click the RSMLT tab.
- 7. Select **Enable**.
- 8. In the **HoldDownTimer** field, type a hold-down timer value.
- 9. In the **HoldUpTimer** field, type a holdup timer value.
- 10. Click Apply.

RSMLT field descriptions

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

Name	Description	
Enable	Enables RSMLT. The default is disabled.	
HoldDownTimer	Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address.	
	The range of this value is from 0 to 3600 seconds. The default is 60.	
	If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.	
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.	

Name	Description	
	If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.	

Enabling RSMLT Edge support

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

The default is disabled.

About this task

RSMLT Edge support configuration applies to both IPv4 and IPv6.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- Click RSMLT.
- 3. Click the Globals tab.
- 4. Select EdgeSupportEnable.
- 5. Click Apply.

Modifying the RSMLT local information

Edit the existing RSMLT configuration for the local node in the cluster.

Procedure

- In the navigation tree, expand the following folders: Configuration > IPv6.
- 2. Click RSMLT.
- 3. Click the Local tab.
- 4. Double-click a cell to change the value.
- 5. Click Apply.

Local field descriptions

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

Name	Description
IfIndex	Shows the route SMLT operation index.
lpv6Addr	Configures the IPv6 address of the RSMLT interface.

Name	Description
Ipv6PrefixLength	Configures the IPv6 prefix length.
Enable	Enables or disables RSMLT. The default is disabled.
HoldDownTimer	Defines how long the recovering/rebooting switch remains in a non-Layer 3 forwarding mode for the peer router MAC address.
	The default is 60.
HoldUpTimer	Defines how long the RSMLT switch maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity.
	The default is 1800.
OperStatus	Displays the RSMLT operating status as either up or down.
Smitld	Specifies the ID range for the SMLT.
VlanId	Configures the VLAN ID.
MacAddr	Configures the MAC address of the VLAN.
Vrfld	Indicates the virtual router ID to which the local RSMLT instance belongs.
VrfName	Indicates the virtual router name to which the local RSMLT instance belongs.

Modifying RSMLT peer information

Edit the existing configuration for the RSMLT peer node in the cluster.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click RSMLT.
- 3. Click the **Peer** tab.
- 4. Double-click a cell to change the value.
- 5. Click Apply.

Peer field descriptions

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

Name	Description
IfIndex	Shows the route SMLT operation index.
lpv6Addr	Configures the IPv6 address of the RSMLT interface.

Name	Description
Ipv6PrefixLength	Configures the IPv6 prefix length.
AdminStatus	Shows the administrative status of RSMLT on the peer.
HoldDownTimer	Defines how long the recovering/rebooting switch remains in a non-Layer 3 forwarding mode for the peer router MAC address.
	The default is 0.
HoldDownTimeRemaining	Indicates the time remaining in the HoldDownTimer.
HoldUpTimer	Defines how long the RSMLT switch maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity.
	The default is 0.
HoldUpTimeRemaining	Indicates the time remaining in the HoldUpTimer.
OperStatus	Displays the RSMLT operating status as either up or down.
Smltld	Specifies the ID range for the SMLT.
VlanId	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.
MacAddr	Configures the MAC address of the VLAN.
Vrfld	Indicates the virtual router ID to which the peer belongs.
VrfName	Indicates the virtual router name to which the peer belongs.

Viewing RSMLT Edge peers

View the RSMLT peers for which the switch acts as a peer forwarder.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click RSMLT.
- 3. Click the **Edge Peers** tab.

Edge Peers field descriptions

Use the data in the following table to use the **Edge Peers** tab.

Name	Description
PeerVlanId	Specifies the ID of the VLAN associated with this entry.
Peerlpv6Address	Specifies the IPv6 address of the peer RSMLT interface.
Peerlpv6PrefixLength	Specifies the peer IPv6 address prefix.
PeerMacAddress	Specifies the peer MAC address.

Chapter 11: Viewing IPv6 connections

This chapter provides procedures you can use to view IPv6 connection information.

You can establish network connectivity with the following protocols:

- Transmission Control Protocol (TCP), for connection-oriented sessions
- User Datagram Protocol (UDP), for connectionless sessions

When you view TCP information you can

- check the health of the connections, from the switch perspective, as they traverse the network
- detect intermittent connectivity
- detect attacks on resources
- · determine which applications are active by checking the port numbers

UDP endpoint information tells you about local and remote UDP activity.

When you view UDP information you can

- determine which applications are active by checking the local and remote port numbers
- identify processes within a UDP session to allow multiplexing of a port mapping for UDP

Viewing IPv6 Connections using CLI

Use the procedures in this section to view IPv6 connections using CLI.

Viewing TCP and UDP information

Perform this procedure to view the TCP and UDP configuration information for IPv6.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display IPv6 TCP connection information:

```
show ipv6 tcp connections [vrf WORD<1-16> | vrfids WORD<0-512>]
```

3. Display IPv6 TCP listener information for the specified IPv6 address:

```
show ipv6 tcp listener [vrf WORD < 1-16 > | vrfids WORD < 0-512 > ]
```

4. Display IPv6 TCP properties

```
show ipv6 tcp properties [vrf WORD<1-16> | vrfids WORD<0-512>]
```

5. Display IPv6 TCP statistics

```
show ipv6 tcp statistics [vrf WORD<1-16> | vrfids WORD<0-512>]
```

6. Display IPv6 UDP information:

show ipv6 udp endpoints [vrf WORD<1-16> | vrfids WORD<0-512>]

Example

```
Switch:1>show ipv6 tcp connections

TCP connection table info - GlobalRouter

LOCALPORT LOCALADDR REMOTEPORT REMOTEADDR STATE

21 0:0:0:0:0:0:0:0:0 0 0:0:0:0:0:0:0 listen
23 0:0:0:0:0:0:0:0:0 0 0:0:0:0:0:0:0 listen
```

Switch:1>show ipv6 tcp listener

TCP listener table info - GlobalRouter

LOCALPORT LOCALADDR

21 0:0:0:0:0:0:0:0
23 0:0:0:0:0:0:0:0
80 0:0:0:0:0:0:0:0:0
443 0:0:0:0:0:0:0:0:0
513 0:0:0:0:0:0:0:0:0

Switch:1#show ipv6 tcp properties

TCP Global Properties - GlobalRouter

RtoAlgorithm constant
RtoMin 5002 milliseconds
RtoMax 60128 milliseconds
MaxConn 127

Switch:1#show ipv6 tcp statistics

TCP Global Statistics - GlobalRouter

ActiveOpens: 0
PassiveOpens: 5
AttemptFails: 0
EstabResets: 0
CurrEstab: 0
InSegs: 0
OutSegs: 0
RetransSegs: 0
OutRsts: 0
OutRsts: 0

HCInSegs: HCOutSegs:	0 0	
Switch:1>sho	ow ipv6 udp endpoints	
	UDP endpoint	table info - GlobalRouter
LOCALPORT REMOTEPORT PROCESS	LOCALADDR REMOTEADDR	INSTANCE
69 0 161	0:0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0:0	1220866096 0 1220867644 0

Viewing IPv6 Connections using EDM

Use the procedures in this section to view IPv6 connections using EDM.

Viewing TCP global information

View TCP and UDP information to view the current configuration.

Before you begin

Change the VRF instance as required to view TCP global information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

The fields on the TCP global tab provide information about the handshake (SYN) configuration and the maximum number of TCP connections you can create on your system.

When you initiate a TCP connection, both end points send handshake information to create the channel.

The retransmission algorithm and fields display the configured timeout value and minimum and maximum retransmission times that your system uses to terminate a connection attempt that falls outside your specified parameters.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders
- 2. Click TCP/UDP.
- 3. Click the **TCP Globals** tab.

TCP Global field descriptions

Use the data in the following table to use the **TCP Globals** tab.

Name	Description
RtoAlgorithm	Determines the timeout value used for retransmitting unacknowledged octets.
RtoMin	Displays the minimum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
RtoMax	Displays the maximum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
MaxConn	Displays the maximum connections for the device.

Viewing TCP connections information

View information about TCP connections.

Before you begin

Change the VRF instance as required to view TCP connections information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

Among other things, the fields on the TCP connections tab provide important information about the health of connections that traverse your switch.

In particular, the state column lets you know the state of each TCP connection. Of these, synSent, synReceived, and established indicate whether or not a channel is established and listen indicates when an end system is waiting for a returning handshake (SYN).

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click TCP/UDP.
- 3. Click the TCP Connections tab.

TCP Connections field descriptions

Use the data in the following table to use the **TCP Connections** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
RemAddressType	Displays the type (IPv6 or IPv4) for the remote address of the TCP connection.

Name	Description
RemAddress	Displays the IPv6 address for the remote TCP connection.
RemPort	Displays the remote port number for the TCP connection.
State	Displays an integer that represents the state for the connection:
	• closed
	• listen
	synSent
	synReceived
	established
	• finWait1
	• finWait2
	closeWait
	• lastAck(9)
	• closing
	timeWait
	• deleteTCB
Process	Displays the process ID for the system process associated with the TCP connection.

Viewing TCP listeners information

View TCP listener information.

Before you begin

Change the VRF instance as required to view TCP listeners information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

The TCP listeners table provides a detailed list of systems that are in the listening state.

When a connection is in the listen state an end point system is waiting for a returning handshake (SYN). The normal listening state should be very transient, changing all of the time.

Two or more systems going to a common system in an extended listening state indicates the need for further investigation.

End systems in an extended listening state can indicate a broken TCP connection or a DOS attack on a resource.

This type of DOS attack, known as a SYN attack, results from the transmission of SYNs with no response to return replies.

While many systems can detect a SYN attack, the TCP listener statistics can provide additional forensic information.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. ClickTCP/UDP.
- 3. Click the TCP Listeners tab.

TCP Listeners field descriptions

Use the data in the following table to use the TCP Listeners tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
Process	Displays the process ID for the system process associated with the TCP connection.

Viewing UDP endpoint information

View UDP Endpoints to confirm correct configuration.

Before you begin

Change the VRF instance as required to view UDP endpoint information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

You can use UDP endpoint information to display local and remote UDP activity.

Since UDP is a protocol used to establish connectionless network sessions, you need to monitor local and remote UDP activity and to know which applications are running over UDP.

You can determine which applications are active by checking the port number.

Processes are further identified with a UDP session to allow for the multiplexing of a port mapping for UDP.

Procedure

- 1. In the navigation pane, expand the **Configuration > IPv6** folders.
- 2. Click TCP/UDP.
- 3. Click the **UDP Endpoints** tab.

UDP Endpoints field descriptions

Use the data in the following table to use the **UDP Endpoints** tab.

Name	Description
LocalAddressType	Displays the local address type (IPv6 or IPv4).
LocalAddress	Displays the local IPv6 address.
LocalPort	Displays the local port number.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IPv6 address.
RemotePort	Displays the remote port number.
Instance	Distinguishes between multiple processes connected to the UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 12: IPv6 Alternative Routes

This chapter provides concepts and procedures to complete IPv6 alternative routes configuration using the CLI.

Alternative routes

Table 22: Alternative routes product support

Feature	Product	Release introduced
Alternative routes for IPv4	VSP 4450 Series	VSP 4000 4.0
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring IPv4 Routing for	VSP 7200 Series	VOSS 4.2.1
<u>VOSS</u> .	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
Alternative routes for IPv6	VSP 4450 Series	VOSS 5.1
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring IPv6 Routing for	VSP 7200 Series	VOSS 5.1
VOSS.	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 5.1
	VSP 8400 Series	VOSS 5.1
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported

To avoid traffic interruption, you can globally enable the alternative routes feature so the router can use the next-best route, also known as an alternative route, if the best route becomes unavailable.

Routers learn routes to a destination through routing protocols. Routers maintain a routing table of the learned alternative routes sorted in order by route preference, route costs, and route sources. The first route on the list is the best route and the route that the router prefers to use.

The alternative route concept also applies between routing protocols. For example, if an OSPFv3 route becomes unavailable and an alternative RIPng route is available, the system activates the RIPng route without waiting for the update interval to expire.

Enable IPv6 Alternative Routes

Use this procedure to enable IPv6 alternative routes and view the configuration on the switch.

Procedure

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

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

2. Enable IPv6 alternative routes:

```
ipv6 alternative-route
```



IPv6 alternative routes are enabled by default.

3. Verify the configuration of the IPv6 alternative route:

```
show ipv6 global [vrf WORD < 1-16 > | vrfids WORD < 0-512 >] show ipv6 route alternative [vrf WORD < 1-16 > | vrfids WORD < 0-512 >]
```

Example:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf globalRouter
Switch:1(router-vrf)#ipv6 alternative-route
```

```
Switch:1#show ipv6 global

IPv6 Global Information - GlobalRouter

forwarding
default-hop-cnt
icmp-error-interval
icmp-error-quota
icmp-unreach-msg
icmp-addr-unreach-msg
icmp-port-unreach-msg
icmp-echo-multicast-request
static-route-admin-status
alternative-route
ecmp
ecmp-max-path

IPv6 Global Information - GlobalRouter

: enable
: 64

1000
: 64

1000
: 50

1 icmp-error-quota
: 50
: enable
: ena
```

	IPv6 Routing Table Information -	GlobalRouter					
Destination Address/PrefixLen	NEXT HOP	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2910:0:0:1:0:0:0:0/64	fe80:0:0:0:b2ad:aaff:fe42:dd00	V-3	OSPF	2	0	В	20
		V-1001	LOCAL	_	0	В	0
	0:0:0:0:0:0:0:0	T-10	BGP	1	0	A	45
3000:0:0:1:0:0:0:0/64	0:0:0:0:0:0:0:0	V-3	LOCAL	1	0	В	0
4001:0:0:1:0:0:0:0/64	0:0:0:0:0:0:0:0	T-10	LOCAL	1	0	В	0
		T-10	BGP	1	0	В	45
	fe80:0:0:0:b2ad:aaff:fe42:dd00	V-3	OSPF	2	0	A	120
5910:0:0:2:0:0:0:0/64	0:0:0:0:0:0:0:0	T-10	BGP	1	0	В	45
5910:0:0:2:0:0:0/64	fe80:0:0:0:b2ad:aaff:fe42:dd00	V-3	OSPF	2	0	A	120

Chapter 13: IPv6 configuration examples

The following sections show configuration examples for IPv6 deployment options.

IPv6 tunnels

This section shows examples of manually configured tunnels between brouter ports and VLANs.

Between brouter ports

The following figure shows the tunnel configuration between brouter ports.

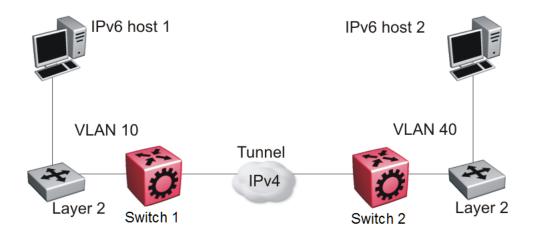


Figure 12: Tunnel configuration between brouter ports

You must configure static routes, RIP, or OSPF on both the source (Switch 1) and destination (Switch 2) IPv4 interfaces to communicate on the IPv4 network. You must configure IPv4 addresses on the source and destination.

Configuring Switch 1

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 10 type port-mstprstp 0
vlan mlt 10 4
vlan members 10 1/1 portmember
interface vlan 10
ipv6 interface
ipv6 interface enable
```

```
ipv6 interface address 4000:0:0:0:0:0:0:1/64 exit.
```

Create an IPv4 brouter port and enable OSPF on the port.

```
interface GigabitEthernet 1/30
brouter port 1/30 vlan 1000 subnet 172.21.80.1/255.255.255.0 mac-offset
6
```

Create the tunnel from the source to the destination.

```
ipv6 tunnel 1 source 172.21.80.1 address 2500:0000:0000:0000:0000:0000:0000:0001/64
destination 192.168.20.1
```

Configure a static route on the source.

```
ipv6 route 4000:0:0:0:0:0:0:2/64 cost 1 tunnel 1
```

Optionally, you can create an OSPFv3 interface through the tunnel.

```
router ospf ipv6-enable
router ospf
ipv6 tunnel 1 area 0.0.0.0
ipv6 tunnel 1 enable
exit
```

Configuring Switch 2

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 40 type port-mstprstp 0
vlan mlt 40 4
vlan members 40 1/2 portmember
interface vlan 40
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000:0:0:0:0:0:2/64
exit
```

Create an IPv4 brouter port and enable OSPF on the port.

```
interface GigabitEthernet 1/30
brouter port 1/30 vlan 2000 subnet 192.168.20.1/255.255.255.0 mac-offset 6
```

Create the tunnel from the destination to the source.

```
ipv6 tunnel 1 source 192.168.20.1 address 2500:0000:0000:0000:0000:0000:0000:0002/64
destination 172.21.80.1
```

Configure a static route on the destination.

```
ipv6 route 4000:0:0:0:0:0:0:1/64 cost 1 tunnel 1
```

Optionally, you can create an OSPFv3 interface through the tunnel.

```
router ospf ipv6-enable
router ospf
ipv6 tunnel 1 area 0.0.0.0
ipv6 tunnel 1 enable
exit
```

Between VLANs

The following figure shows the tunnel configuration between VLANs.

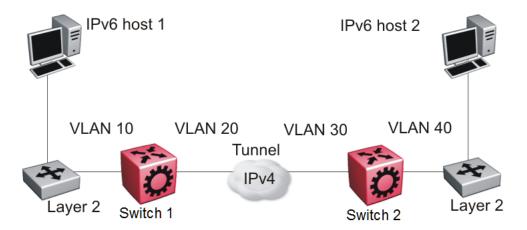


Figure 13: Tunnel configuration between VLANs

You must configure static routes, and either RIP or OSPF on both the source (Switch 1) and destination (Switch 2) IPv4 interfaces to communicate on the IPv4 network. You must configure IPv4 addresses on the VLANs.

Configuring Switch 1

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 10 type port-mstprstp 0
vlan mlt 10 4
vlan members 10 1/1 portmember
interface vlan 10
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000:0:0:0:0:0:1/64
exit
```

Create an IPv4 VLAN, add ports to the VLAN, and enable OSPF on the VLAN.

```
vlan create 20 type port-mstprstp 0
vlan mlt 20 4
vlan members 20 1/30 portmember
interface vlan 20
ip address 172.21.80.1 255.0.0.0
ip ospf enable
exit
```

Create the tunnel from the source to the destination.

```
ipv6 tunnel
1 source 172.21.80.1 address 2500:0000:0000:0000:0000:0000:0000:0001/64
destination 192.168.20.1
```

Configuring Switch 2

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 40 type port-mstprstp 0
vlan mlt 40 4
vlan members 40 1/2 portmember
interface vlan 40
ipv6 interface
ipv6 interface enable
```

ipv6 interface address 4000:0:0:0:0:0:0:2/64 exit.

Create an IPv4 VLAN, add ports to the VLAN, and enable OSPF on the VLAN.

```
vlan create 30 type port-mstprstp 0 vlan mlt 30 4 vlan members 30 1/30 portmember interface vlan 30 ip address 192.168.20.1 255.0.0.0 ip ospf enable exit
```

Create the tunnel from the destination to the source.

ipv6 tunnel 1 source 192.168.20.1 address 2500:0000:0000:0000:0000:0000:0000:0002/64
destination 172.21.80.1

Verification

Use the following show command to verify tunnel creation on the source device:

```
Switch:1(config) #show ipv6 tunnel 1 detail
______
                            _____
                  Tunnel Interface Information
      LOCAL ADDRESS REMOTE ADDRESS OPER STATUS TYPE
TD
      172.21.80.1 192.168.20.1 active manual
1 out of 1 Total number of entries displayed.
                    Address Information
______
TPV6
                            TYPE
                                 ORIGIN STATUS
ADDRESS
                           UNICAST MANUAL PREFERRED
2500:0:0:0:0:0:1
fe80:0:0:0:0:0:ac15:5001
                            UNICAST LINKLAYER PREFERRED
2 out of 2 Total number of entries displayed.
```

Use the following show command to verify tunnel creation on the destination device:

```
Switch:1(config) #show ipv6 tunnel 1 detail
______
          Tunnel Interface Information
______
 LOCAL ADDRESS REMOTE ADDRESS OPER STATUS TYPE
______
    192.168.20.1
          172.21.80.1
                 active
                     manual
1 out of 1 Total number of entries displayed.
______
            Address Information
______
IPV6
                TYPE ORIGIN STATUS
ADDRESS
```

```
2500:0:0:0:0:0:2 UNICAST MANUAL PREFERRED fe80:0:0:0:0:0:c0a8:1401 UNICAST LINKLAYER PREFERRED

2 out of 2 Total number of entries displayed.
```

OSPFv3

This section shows an example of OSPFv3 configuration. The following figure shows the network.

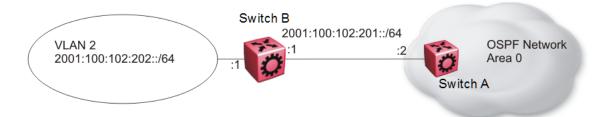


Figure 14: OSPFv3 configuration

To complete the configuration, you must perform the following actions:

- Configure an IPv6 VLAN (VLAN 2) with port member 1/1.
- Configure an IPv6 brouter port (1/2).
- Use IPv6 address 2001:100:102::/64.

Configure VLAN 2 and add port members.

```
vlan create 2 type port-mstprstp 0
vlan mlt 2 4
vlan members 2 1/1 portmember
interface vlan 2
ipv6 interface
ipv6 interface enable
ipv6 interface address 2001:100:102:202:0:0:1/64
exit
```

Enable OSPFv3 on VLAN 2.

```
# IPV6 OSPF VLAN CONFIGURATION
interface vlan 2
ipv6 ospf area 0.0.0.0
ipv6 ospf poll-interval 0
ipv6 ospf enable
exit
```

Create brouter port 1/2 with IPv6 and OSPFv3.

```
interface gigabitethernet 1/2
ipv6 interface vlan 3999
ipv6 interface enable
ipv6 interface address 2001:100:102:201:0:0:0:1/64
ipv6 ospf area 0.0.0.0
```

ipv6 ospf enable
exit

Verification

The following example shows the Global Router example for OSPFv3 Area configuration:

```
Switch:1#show ipv6 ospf area

OSPF Area - GlobalRouter

AREA_ID STUB_AREA NSSA IMPORT_SUM TRANS_ROLE TRANS_STATE

0.0.0.0 false false true always disabled
STUB_METRIC_TYPE_SPF_RUNS_BDR_RTR_CNT_ASBDR_RTR_CNT_LSA_CNT_LSACK_SUM

10 ospfv3Metric 0 0 0 0 0 0
```

The following example shows the VRF example for OSPFv3 Area configuration:

Switch:1#show ipv6 ospf area vrf vrf1								
OSPF Area - VRF vrf1								
AREA_ID	STUB_AREA	NSSA I	IMPORT_SUM T	RANS_ROLE T	TRANS_STA	TE		
0.0.0.0 1.1.1.1 STUB_METRIC		false t		andidate d	disabled	LSA_CNT	LSACK_	SUM
10	ospfV3Metric ospfV3Metric	3 3	0	0 1		0	0 0	

Switch:1#show ipv6 interface vlan 2							
Vlan Ipv6 Interfa	==== ace		======				======
IF VLAN PHYSICAL ADMIN OPER TYPE MTU INDX INDX ADDRESS STATE STATE		ABLE	SMIT	MCAST STATU		RPC	RPC MODE
2070 2 00:24:7f: enable up ETHER 1500 al:7a:06	64	30000	1000	disab	le disak	ole	exist
Vlan Ipv	76 Ad	ldress					
IPV6 ADDRESS	VLA	N-ID	TY1	 PE	ORIGIN	S'	TATUS
					ANUAL INKLAYEF		
1 out of 2 Total Num of Interface Entri 2 out of 5 Total Num of Address Entries			•				

IPv6 alternative routes configuration example

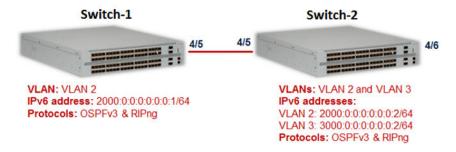
To avoid traffic interruption, you can enable alternative routes globally to replace the best route with the next-best route, if the best route becomes unavailable.

The concept of alternative route applies between routing protocols. For example, if an OSPFv3 route becomes unavailable and an alternative RIPng route is available, the system activates the RIPng route immediately without waiting for an update interval to expire.

By default, the alternative routes feature is globally enabled on the switch.

The following example demonstrates this behavior.

In this example, you configure OSPFv3 and RIPng routes on two switches Switch-1 and Switch-2, as shown in the following figure.



Configuration on Switch-1

VLAN configuration:

On Switch-1, configure VLAN 2 and the IPv6 interface address 2000:0:0:0:0:0:0:0:1/64.

```
Switch1:1:1>enable
Switch1:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config) #vlan create 2 type port-mstprstp 0
Switch1:1(config) #vlan members 2 4/5
Switch1:1(config) #interface vlan 2
Switch1:1(config-if)#ipv6 interface address 2000:0:0:0:0:0:0:1/64
Switch1:1(config-if) #ipv6 interface enable
Switch1:1(config-if)#exit
Switch1:1(config) #show vlan basic
______
                   Vlan Basic
_____
VLAN
                     MSTP
ID NAME TYPE
                     INST ID PROTOCOLID SUBNETADDR
                                                  SUBNETMASK
                                                                VRFID
    Default byPort 0 none VLAN-2 byPort 0 none
                                            N/A
                                 N/A
                                                                0
                                                                0
                                      N/A
All 2 out of 2 Total Num of Vlans displayed
Switch1:1(config) #show vlan members
                           Vlan Port
```

```
______
VLAN PORT ACTIVE STATIC NOT_ALLOW ID MEMBER MEMBER MEMBER MEMBER
  1/1-1/16,1/17/1- 1/1-1/16,1/17/1-
1/17/4,1/18/1- 1/17/4,1/18/1-
1/18/4,2/1-2/16, 1/18/4,2/1-2/16,
2/17/1-2/17/4, 2/17/1-2/17/4,
2/18/1-2/18/4,3/1- 2/18/1-2/18/4,3/1-
     3/6,4/1-4/4,4/6 3/6,4/1-4/4,4/6
     4/5
                           4/5
All 2 out of 2 Total Num of Port Entries displayed
Switch1:1(config) #show ipv6 interface vlan 2
                     Vlan Ipv6 Interface
______
IFINDX VLAN PHYSICAL ADMIN OPER TYPE MTU HOP REACHABLE RETRANSMIT MCAST IPSEC RPC INDX ADDRESS STATE STATE LMT TIME TIME STATUS
                                                                                         RPCMODE
2050 2 b0:ad:aa:4e:59:00 enable up ETHER 1500 64 30000 1000 disable disable disable existonly
______
                        Vlan Ipv6 Address
______

        IPV6 ADDRESS
        VLAN-ID
        TYPE
        ORIGIN
        STATUS

        2000:0:0:0:0:0:0:0:1/64
        V-2
        UNICAST MANUAL PREFERRED

        fe80:0:0:0:0:b2ad:aaff:fe4e:5900/64
        V-2
        UNICAST LINKLAYER PREFERRED

1 out of 1 Total Num of Interface Entries displayed.
2 out of 2 Total Num of Address Entries displayed.
```

Port configuration:

```
Switch1:1(config)#interface gigabitEthernet 4/5
Switch1:1(config-if)#encapsulation dot1q
Switch1:1(config-if)#no shutdown
Switch1:1(config-if)#exit
```

IPv6 global configuration:

```
Switch1:1(config) #ipv6 forwarding

Switch1:1(config) #show ipv6 forwarding

Ipv6 forwarding - GlobalRouter : enable

ecmp
ecmp : disable
ecmp-max-path : 1
```

IPv6 OSPFv3 VLAN configuration:

transit-delay : 1 type : BROADCAST

IPv6 OSPFv3 router configuration:

```
Switch1:1(config-if)#exit
Switch1:1(config) #router ospf ipv6-enable
Switch1:1(config) #show ipv6 ospf
                 OSPFv3 Global Information - GlobalRouter
      router-id
                                   : 170.78.88.0
                                    : ENABLED
       admin-state
       version
       area-bdr-rtr-state
as-bdr-rtr-state
                                    : FALSE
                                    : FALSE
      helper-mode
                                   : ENABLED
       as-scope-lsa-count
                                    : 0
       lsa-checksum
       originate-new-lsas
                                    : 22
                                    : 11
       rx-new-lsas
       ext-lsa-count
```

Switch1:1(config) #show ipv6 ospf neighbor ______ OSPF Neighbor - GlobalRouter ______ IFINDX (VID/BRT) NBRROUTERID NBRIPADDR STATE TTL 2050 (2) 170.78.84.0 fe80:0:0:0:b2ad:aaff:fe4e:5500 Full 31 1 out of 1 Total Num of Neighbor Entries displayed. _____ OSPF Virtual Neighbor - GlobalRouter ______ NBRAREAID NBRROUTERID VIRTINTFID NBRIPV6ADDR STATE ______ 0 out of 0 Total Num of Virtual Neighbor Entries displayed. _____ OSPF NBMA Neighbor - GlobalRouter ------INTERFACE NBRROUTERID NBRIPADDR STATE 0 out of 0 Total Num of NBMA Neighbor Entries displayed. H = Helping a Restarting neighbor Switch1:1(config-if)#exit

IPv6 RIPng configuration on VLAN:

```
Switch1:1(config) #interface vlan 2
Switch1:1(config-if) #ipv6 rip
Switch1:1(config-if) #ipv6 rip enable

Switch1:1(config-if) #show ipv6 rip interface

Total RIPng interfaces: 1

RIPng Interface - GlobalRouter

FINDX COST POISON SEND ADMIN OPER STATUS

STATUS DEFAULT STATUS STATUS

2050 (2 ) 1 disable disable enable enable

1 out of 1 Total Num of RIPng interfaces displayed
```

IPv6 RIPng global router configuration:

```
Switch1:1(config) #router rip ipv6-enable
Switch1:1(config) #show ipv6 rip

Switch1:1(config) #show ipv6 rip

RIPng Global - GlobalRouter

Rip: Enabled
HoldDown Time: 120
Timeout Interval: 180
Update Time: 30
Default Info Metric: 1
Default Info State: Disabled
Default Import Metric: 1
```

Configuration on Switch-2

On Switch-2, configure VLAN 2 and VLAN 3 with the IPv6 interfaces 2000:0:0:0:0:0:0:0:2/64 and 3000:0:0:0:0:0:0:0:2/64 respectively.

VLAN configuration:

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch2:1(config) #vlan create 2 type port-mstprstp 0
Switch2:1(config) #vlan members 2 4/5 portmember
Switch2:1(config) #interface vlan 2

Switch2:1(config-if) #ipv6 interface address 2000:0:0:0:0:0:0:2/64
Switch2:1(config-if) #ipv6 interface enable
Switch2:1(config-if) #ipv6 forwarding
Switch2:1(config-if) #exit

Switch2:1(config) #vlan create 3 type port-mstprstp 0
Switch2:1(config) #vlan members 3 4/6 portmember
Switch2:1(config) #vlan members 3 4/6 portmember
```

```
Switch2:1(config-if) #ipv6 interface address 3000:0:0:0:0:0:0:2/64
Switch2:1(config-if) #ipv6 interface enable
Switch2:1(config-if) #ipv6 forwarding
Switch2:1(config-if) #exit
```

Vlan Basic										
VLAN ID	NAME	TYPE	MSTP INST_ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRFID			
1	Default	byPort	0	none	N/A	N/A	0			
2	VLAN-2	byPort	0	none	N/A	N/A	0			
3	VLAN-3	byPort	0	none	N/A	N/A	0			

All 3 out of 3 Total Num of Vlans displayed

Switch2:1(config) #show vlan members

Vlan	Port
------	------

		Vian FO.	LL	
VLAN ID	PORT MEMBER	ACTIVE MEMBER	STATIC MEMBER	NOT_ALLOW MEMBER
1		1/17/4,1/18/1- 1/18/4,2/1-2/16, 2/17/1-2/17/4,		
2	4/5	4/5		
3	4/6	4/6		

All 3 out of 3 Total Num of Port Entries displayed

Switch2:1(config) #show ipv6 interface vlan

		Vlan	Ipv6 In	terface	9								
IFINDX INDX	VLAN		ADMIN STATE	OPER STATE	TYPE	MTU	HOP LMT		RETRANSMIT TIME	MCAST STATUS	IPSEC	RPC	RPCMODE
2050 2051		b0:ad:aa:4e:55:00 b0:ad:aa:4e:55:01						30000 30000	1000 1000				existonly existonly
======			Vlan I	===== pv6 Ado	dress			======			======	======	======
IPV6 ADDRESS				VLAI	N-ID	Т	YPE	ORIGIN	STATUS		======	======	======
fe80:0:0:0:b2ad:aaff:fe4e:5500/64 3000:0:0:0:0:0:2/64				V-2 V-2 V-3 V-3		U:	NICAS NICAS	r manual	PREFERRED R PREFERRED PREFERRED R PREFERRED				
		Total Num of Interi Total Num of Addres											

Port configuration:

```
Switch2:1(config) #interface GigabitEthernet 4/5
Switch2:1(config) #encapsulation dot1q
Switch2:1(config) #no shutdown
```

```
Switch2:1(config) #interface GigabitEthernet 4/6
Switch2:1(config) #encapsulation dot1q
Switch2:1(config) #no shutdown
```

IPv6 global configuration:

```
Switch1:1(config)#ipv6 forwarding
Switch1:1(config)#show ipv6 forwarding
Ipv6 forwarding - GlobalRouter : enable
ecmp
ecmp
ecmp-max-path
: 1
```

IPv6 OSPFv3 VLAN configuration:

```
Switch2:1(config)#interface vlan 2
Switch2:1(config-if)#ipv6 ospf area 0.0.0.0
Switch2:1(config-if)#ipv6 ospf enable

Switch2:1(config)#interface vlan 3
Switch2:1(config-if)#ipv6 ospf area 0.0.0.0
Switch2:1(config-if)#ipv6 ospf enable

Switch2:1(config-if)#show ipv6 ospf
```

------OSPFv3 Global Information - GlobalRouter

router-id : 170.78.84.0

admin-state : ENABLED

version : 3

area-bdr-rtr-state : FALSE

as-bdr-rtr-state : FALSE

helper-mode : ENABLED

as-scope-lsa-count : 0

lsa-checksum : 0

originate-new-lsas : 56

rx-new-lsas : 62

ext-lsa-count : 0

Switch2:1(config-if)#show ipv6 ospf interface
Total ospf areas: 1
Total ospf interfaces: 2

OSPF Interface - GlobalRouter

```
IFINDX(VID/BRT) AREAID ADM IFSTATE METRIC PRI DR/BDR IFTYPE

2050 (2 ) 0.0.0.0 ena BDR 1 1 170.78.88.0 BROADCAST 170.78.84.0
2051 (3 ) 0.0.0.0 ena DR 1 1 170.78.84.0 BROADCAST 0.0.0.0
```

2 out of 2 Total Num of ospf interfaces displayed

Total ospf virtual interfaces: 0

OSPF Virtual Interface - GlobalRouter

```
______
AREAID
       NBRIPADDR
                STATE
0 out of 0 Total Num of ospf virtual interfaces displayed
Switch2:1(config-if) #show ipv6 ospf neighbor
______
              OSPF Neighbor - GlobalRouter
_____
IFINDX (VID/BRT) NBRROUTERID NBRIPADDR
                                 STATE TTL
2050 (2) 170.78.88.0 fe80:0:0:0:b2ad:aaff:fe4e:5900 Full 30
1 out of 1 Total Num of Neighbor Entries displayed.
_____
            OSPF Virtual Neighbor - GlobalRouter
_____
       NBRROUTERID VIRTINTFID NBRIPV6ADDR
NBRAREATD
                                 STATE
0 out of 0 Total Num of Virtual Neighbor Entries displayed.
______
             OSPF NBMA Neighbor - GlobalRouter
______
INTERFACE NBRROUTERID NBRIPADDR
                                 STATE
0 out of 0 Total Num of NBMA Neighbor Entries displayed.
H = Helping a Restarting neighbor
```

IPv6 OSPFv3 global router configuration:

```
Switch2:1(config-if)#exit
Switch2:1(config) #router ospf ipv6-enable
Switch1:1(config) #show ipv6 ospf
______
           OSPFv3 Global Information - GlobalRouter
______
                          : 170.78.88.0
     router-id
                          : ENABLED
     admin-state
     version
                          : 3
     area-bdr-rtr-state
                          : FALSE
     as-bdr-rtr-state
                          : FALSE
     helper-mode
                          : ENABLED
     as-scope-lsa-count
                          : 0
     lsa-checksum
                          : 0
     originate-new-lsas
```

```
rx-new-lsas : 11
ext-lsa-count : 0
```

IPv6 RIPng configuration:

```
Switch2:1(config) #interface vlan 2
Switch2:1(config-if)#ipv6 rip
Switch2:1(config-if) #ipv6 rip enable
Switch2:1(config-if)#exit
Switch2:1(config) #interface vlan 3
Switch2:1(config-if)#ipv6 rip
Switch2:1(config-if) #ipv6 rip enable
Switch2:1(config-if)#exit
Switch2:1(config)#
Switch2:1(config) #show ipv6 rip interface
Total RIPng interfaces: 2
______
            RIPng Interface - GlobalRouter
IFINDX COST POISON SEND ADMIN OPER STATUS DEFAULT STATUS STATUS
2050 (2 ) 1 disable disable enable enable 2051 (3 ) 1 disable disable enable
2 out of 2 Total Num of RIPng interfaces displayed
```

IPv6 RIPng global router configuration:

Viewing route and alternative route configuration on the switches

On Switch-1 and Switch-2, the route 3000:0:0:0:0:0:0:0:0:2/64 is learned using the protocols RIPng and OSPFv3. The OSPFv3 route is learned as the best route because of its route preference value of 20. The RIPng route is added as alternative route as it has the route preference 100, which is greater than the OSPFv3 route preference of 20. On Switch-2, the route 3000:0:0:0:0:0:0:0:0:0:0:0:0:0 is a local route.

Viewing route and alternative route configuration on Switch-1:

```
Switch1:1(config) #show ipv6 route alternative

IPv6 Routing Table Information - GlobalRouter
```

Destination Address/PrefixLen	NEXT HOP V	ID/DID/IID F					
	0:0:0:0:0:0:0:0:0						
	fe80:0:0:0:b2ad:aaff:fe4e:550 fe80:0:0:0:b2ad:aaff:fe4e:550		RIP	2 2		0 В 0 А	1
4 out of 4 Total Num of Route	Entries displayed.						
TYPE Legend: A=Alternative Route, B=Best R	oute, E=Ecmp Route						
Switch1:1(config)#show ipv6 r	oute						
Switch1:1(config)#show ipv6 r	oute						====
======================================	 Routing Table Information - Gl			=====		====	====
IPv6	Routing Table Information - Gl						
Destination Address/PrefixLen	Routing Table Information - Gl	VID/BID/TID	PROTO	COST	AGE	TYPE	PREE
IPv6 Destination Address/PrefixLen 2000:0:0:0:0:0:0:1/64	Routing Table Information - Gl	VID/BID/TID	PROTO LOCAL	COST 1	AGE 0	TYPE B	PREE
IPv6 Destination Address/PrefixLen 2000:0:0:0:0:0:0:1/64	Routing Table Information - G1	VID/BID/TID	PROTO LOCAL	COST 1	AGE 0	TYPE B	PREE

Viewing route and alternative route configuration on Switch-2:

IPv	6 Routing Table Inf		lobalRoute	er					
Destination Address/PrefixLen			VID/BID,	/TID	PROTO	COS	T AGE	TYPE	PREI
2000:0:0:0:0:0:0:2/64 3000:0:0:0:0:0:0:2/64 3000:0:0:0:0:0:0:2/64	0:0:0:0:0:0:0:0 fe80:0:0:0:b2ad:aa	ff:fe4e:5500	V-2 V-2		LOCAL LOCAL	1 2	0	B B	0 20 100
4 out of 4 Total Num of Route	Entries displayed.								
TYPE Legend: A=Alternative Route, B=Best Ro	ute, E=Ecmp Route								
Switch2:1#show ipv6 route									
	Routing Table Info	rmation - Glo			=====	=====	===		
Destination Address/PrefixLen	NEXT HOP	VID/BID/TID			AGE				
2000:0:0:0:0:0:0:2/64 3000:0:0:0:0:0:0:2/64	0:0:0:0:0:0:0:0	V-2 V-3	LOCAL :	L L	0	B B			

Changing the route preference on Switch-1

On the switch, default preferences are assigned to all standard routing protocols. You can modify the global preference for a protocol to give it a higher or lower priority than other protocols. When you change the preference for a static route, if all best routes remain best routes, only the local route tables change. However, if changing the protocol preference causes best routes to no longer be best routes, neighboring route tables can be affected.

In the following example scenario, you configure a different routing preference for the RIPng protocol on Switch-1 and observe the learning of best and alternative routes. The existing route preference for RIPng is 100.

```
Switch1:1#show ipv6 route alternative

IPv6 Routing Table Information - GlobalRouter

Destination Address/PrefixLen NEXT HOP VID/BID/TID PROTO COST AGE TYPE PREF
```

Configure a different route preference for the RIPng protocol, for example, 19:

```
Switch1:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config)#ipv6 route preference protocol ripng 19
Switch1:1(config)#exit
```

Verify the route preference configuration:

```
Switch1:1#show ipv6 route preference

IPv6 Route Preference - GlobalRouter

PROTOCOL DEFAULT CONFIG

LOCAL 0 0 0
STATIC 5 5
SPBM_L1 7 7
OSPFV3_INTRA 20 20
OSFFV3_INTRA 20 20
OSFFV3_INTER 25 25
EBGP 45 45
RIPNG 100 19
OSPFV3_E1 120 120
OSPFV3_E2 125 125
IBGP 175 175
```

View the updated route preference (for RIPng) on Switch-1. The RIPng route is now learnt as the best route as it has lesser value of route preference (19) than that of OSPFv3 (20), as shown below.

Destination Address/PrefixLen	NEXT HOP	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2000:0:0:0:0:0:0:1/64 3000:0:0:0:0:0:0:2/64 3000:0:0:0:0:0:0:2/64	0:0:0:0:0:0:0:0:0 fe80:0:0:0:b2ad:aaff:fe4e:5500 fe80:0:0:0:b2ad:aaff:fe4e:5500	V-2 V-2 V-2	LOCAL RIP OSPF	1 2 2	0 0 0	B B A	0 19 20
4 out of 4 Total Num of Route Entries displayed.							
TYPE Legend: A-Alternative Route, B=Best Route, E=Ecmp Route							

Disable alternative route learning on Switch-1

The following example demonstrates disabling alternative route learning on Switch-1.

View the alternative routes on Switch-1.

```
IPv6 Routing Table Information - GlobalRouter

Destination Address/PrefixLen NEXT HOP VID/BID/TID PROTO COST AGE TYPE PREF

2000:0:0:0:0:0:1/64 0:0:0:0:0:0:0

3000:0:0:0:0:0:0:2/64 fe80:0:0:0:b2ad:aaff:fe4e:5500 V-2 OSFF 2 0 B 20

3000:0:0:0:0:0:0:2/64 fe80:0:0:0:b2ad:aaff:fe4e:5500 V-2 RIP 2 0 A 100

4 out of 4 Total Num of Route Entries displayed.

TYPE Legend:
A-Alternative Route, B-Best Route, E-Ecmp Route
```

Disable IPv6 alternative routes on Switch-1.

```
Switch1:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config) #no ipv6 alternative-route
Switch1:1(config) #exit
```

Verify that alternative route learning is disabled.

```
Switchl:1#show ipv6 global
forwarding : enable
default-hop-cnt : 64
number-of-interfaces : 1
icmp-error-interval : 1000
icmp-error-quota : 50
icmp-unreach-msg : disable
icmp-echo-multicast-request : enable
static-route-admin-status : enable
alternative-route : disable
ecmp
ecmp-max-path
source-route : disable
```

```
IPv6 Routing Table Information - GlobalRouter

Destination Address/PrefixLen NEXT HOP

VID/BID/TID PROTO COST AGE TYPE PREF

2000:0:0:0:0:0:0:1/64 0:0:0:0:0:0:0:0 V-2 LOCAL 1 0 B 0
3000:0:0:0:0:0:0:0/64 fe80:0:0:0:b2ad:aaff:fe4e:5500 V-2 OSPF 2 0 B 20

3 out of 3 Total Num of Route Entries displayed.

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route
```

Note that the alternative route (RIPng) is not learnt.

```
Switch1:1(config) #show ipv6 route alternative
```

IPv6 Routing Table Information - GlobalRouter							
Destination Address/PrefixLen	NEXT HOP	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2000:0:0:0:0:0:1/64 3000:0:0:0:0:0:0:2/64	0:0:0:0:0:0:0:0 fe80:0:0:0:b2ad:aaff:fe4e:5500	V-2 V-2	LOCAL OSPF			В В	0 20
3 out of 3 Total Num of Route Entries displayed.							
TYPE Legend: A-Alternative Route, B=Best Route, E=Ecmp Route							

Chapter 14: VRF Lite

Table 23: IPv6 virtualization product support

Feature	Product	Release introduced
IPv6 Virtualization for the following	VSP 4450 Series	VOSS 7.0
features and functions:	VSP 4900 Series	VOSS 8.1
IPv6 Interfaces and IPv6 Static	VSP 7200 Series	VOSS 7.0
Routes in VRFs and Layer 3 VSNs	VSP 7400 Series	VOSS 8.0
ECMP and Alternative route	VSP 8200 Series	VOSS 7.0
Route redistribution for static	VSP 8400 Series	VOSS 7.0
and direct routes • VRRPv3 for IPv6	VSP 8600 Series	Not Supported
DHCP RelayIPv6 Reverse Path ForwardingICMP Ping and Traceroute	XA1400 Series	Not Supported
For configuration details, see Configuring IPv6 Routing for VOSS.		
IPv6 Virtualization for the following	VSP 4450 Series	VOSS 8.0
features and functions:	VSP 4900 Series	VOSS 8.1
Open Shortest Path First for IPv6 (OSPFv3)	VSP 7200 Series	VOSS 8.0
IPv6 Border Gateway Protocol	VSP 7400 Series	VOSS 8.0
(IPv6 BGP)	VSP 8200 Series	VOSS 8.0
IPv6 route redistribution	VSP 8400 Series	VOSS 8.0
enhancements	VSP 8600 Series	Not Supported
For configuration details, see Configuring IPv6 Routing for VOSS.	XA1400 Series	Not Supported
iBGP over user-created VRFs	VSP 4450 Series	VOSS 8.1
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring BGP Services for	VSP 7200 Series	VOSS 8.1
<u>VOSS</u> .	VSP 7400 Series	VOSS 8.1
		Toble continues

Feature	Product	Release introduced
	VSP 8200 Series	VOSS 8.1
	VSP 8400 Series	VOSS 8.1
	VSP 8600 Series	Not Supported
	XA1400 Series	Not Supported

VRF Lite provides secure customer data isolation.

VRF Lite Fundamentals

The switch supports what is termed as VRF Lite. Lite conveys the fact that the switch does not use Multiprotocol Label Switching (MPLS) for VRF; VRF Lite is a device virtualization feature, not a network-wide virtualization feature.

Use VRF Lite to offer networking capabilities and traffic isolation to customers that operate over the same node (router). Each virtual router emulates the behavior of a dedicated hardware router; the network treats each virtual router as a separate physical router. In effect, you can perform the functions of many routers using a single platform that runs VRF Lite. The result is a substantial reduction in the cost associated with providing routing and traffic isolation for multiple clients.

With multicast virtualization for IPv4, the switch can function as multiple virtual multicast routers.

The following figure shows one platform acting as multiple virtual routers, each serving a different customer network.

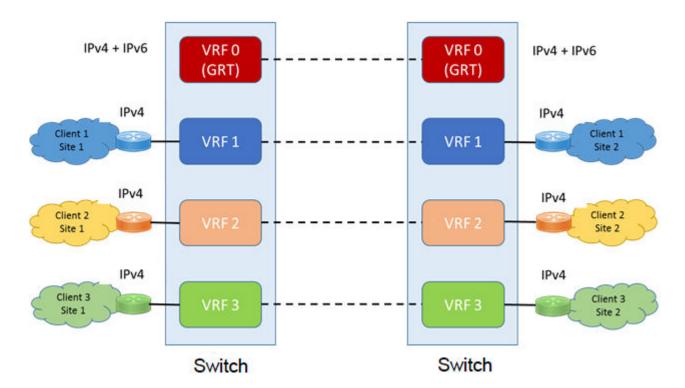


Figure 15: Multiple virtual routers in one system

A switch can support many virtual routers. Each virtual router instance is called a VRF instance. A VRF represents a single instance of a virtual router. Each instance maintains its own routing table. The term Multiple Virtual Router (MVR) is sometimes used to represent a router that contains many VRF instances.

The IPv6 Virtualization functionality adds IPv6 support on VRFs and Layer 3 VSNs. Each VRF instance has its own IPv6 interfaces, IPv6 address space, IPv6 routing table, and IPv6 global parameters. For more information on Layer 3 VSN, see Configuring Fabric Layer 3 Services for VOSS.

The Global Router, VRF 0, is the first instance of the router. When the system starts, it creates VRF 0 by default. VRF 0 provides all non-virtual and traditional routing services. You cannot delete this instance. You can create and configure other VRF instances, if required.

VRF 0 is the only VRF that you can log into through CLI. CLI requires you to specify the VRF when you enter commands.

You can associate one VRF instance with many IPv4 or IPv6 interfaces. These interfaces are unique for each VRF instance. An interface is an entity with an IPv4 or IPv6 address that has the following characteristics:

- A unique association with a VLAN.
- A unique association with a brouter, if not associated with a VLAN
- A unique association with a circuit

A VLAN can only be associated with a single VRF instance.

Note:

- You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IPv4 or IPv6 address. You must first associate the port and VRF instance and then you can configure the IPv4 or IPv6 address.
- Use command boot config flag vrf-scaling to increase total VRFs. You must have a premier license to increase the total VRF count on the switch. For more information on route scaling, see <u>Release Notes for VOSS</u>.

Management VRF

The following sections detail Management VRF features.

Management Port

The management port is a 10/100/1000 Mb/s Ethernet port that you can use for an out-of-band management connection to the switch.



Not all hardware platforms provide a dedicated, physical management interface. For more information, see your hardware documentation.

Management Router VRF

The switch has a separate VRF called Management Router (MgmtRouter) reserved for the management port and the Virtual Management IP address. The configured IP subnet has to be globally unique because the management protocols, for example, SNMP, Telnet, and FTP, can go through in-band or out-of-band ports. The VRF ID for the Management Router is 512.

The switch never switches or routes transit packets between the Management Router VRF port and the Global Router VRF, or between the Management Router VRF and other VRF ports.

The switch honors the VRF of the ingress packet; however, in no circumstance does the switch allow routing between the Management VRF and Global Router VRF. The switch does not support the configuration if you have an out-of-band management network with access to the same networks present in the GRT routing table.



IPv6 is not supported on MgmtRouter.

Non Virtualized Client Management applications

Ensure that you do not define a default route in the Management Router VRF. A route used for non-virtualized client management applications, such as Telnet, Secure Shell (SSH), and FTP, originating from the switch, will always match a default route defined in the Management Router VRF.

If you want out-of-band management, it is recommended that you define a specific static route in the Management Router VRF to the IP subnet where your management application resides.

When you specify a static route in the Management Router VRF, it enables the client management applications originating from the switch to perform out-of-band management without affecting inband management. This enables in-band management applications to operate in the Global Router VRF.

Non-virtualized client management applications originating from the switch, such as Telnet, SSH, and FTP, follow the behavior listed below:

- 1. Look at the Management Router VRF route table.
- 2. If no route is found, the applications will proceed to look in the Global Router VRF table.

Non-virtualized client management applications include:

- DNS
- FTP client with the copy command
- NTP
- rlogin
- · RADIUS authentication and accounting
- SSH
- SNMP clients in the form of traps
- SYSLOG
- TACACS+
- Telnet
- TFTP client

For management applications that originate outside the switch, the initial incoming packets establish a VRF context that limits the return path to the same VRF context.

Virtualized Management Applications

Virtualized management applications, such as ping and traceroute, operate using the specified VRF context. To operate ping or traceroute you must specify the desired VRF context. If not specified, ping defaults to the Global Router VRF. For example, if you want to ping a device through the out-of-band management port you must select the Management Router VRF. For example, if you want to ping a device through the out-of-band management port you must select the Management Router VRF.



Ipv6 is not supported on MgmtRouter.

Switch:1(config) #ping 192.0.2.1 vrf MgmtRouter 192.0.2.1 is alive

Ping test for IPv6:

Switch:1(config) #ping 2001:db8::1 vrf vrfRED 2001:db8::1 is alive

Traceroute test for IPv4:

Switch:1#traceroute 192.0.2.1 vrf MgmtRouter

Traceroute test for IPv6:

Switch:1#traceroute 2001:db8::1 vrf vrfRED

VRF Lite configuration rules

You must select the VRF for global IPv4 or IPv6 options before entering commands.

Not all Global Router parameters are configurable on other VRF instances.

For instructions about how to configure a VRF instance, see the following paragraphs.

Layer 1 and Layer 2 information (including VLAN information) is global and is not maintained for each VRF instance. However, you can associate a set of VLANs with a VRF instance.

One VLAN cannot belong to more than one VRF instance at one time. When you create a VLAN, more than one physical port can belong to it. You can associate a VRF instance with more than one IPv4 or IPv6 interface (a physical Ethernet port or a VLAN).

Perform physical port assignment at the VLAN and brouter port level. A VRF instance inherits all the ports assigned to its VLANs and brouter ports. You cannot directly assign a physical port to a VRF instance, but it is implicitly assigned when you associate the VRF with VLANs or brouter ports.

For IPv4, after you configure interVRF route redistribution between two VRF instances, avoid overlapping IP addresses between these two VRF instances.

When you configure VRF Lite, remember the following rules:

- You can connect two VRFs from the same system with an external cable.
- An IPv4 or IPv6 routable VLAN can become a member of a VRF.
- An IPv4 or IPv6 interface can belong to only one VRF.
- A VRF can exist even if no interfaces are assigned to it.
- Routing policies apply to VRFs on an individual basis.
- Multiple VRFs on the same node can function in different autonomous systems.

Following rules apply to IPv4 interfaces specifically:

- If you configure an IPv4 interface without specifying the VRF instance, it is mapped to VRF 0 by default.
- VRF Lite supports SMLT and RSMLT
- VRF Lite supports RIP in and out policies
- VRF Lite supports OSPF in and out (accept and redistribute) policies

- Before you delete a VRF instance, disable OSPF. Deleting a VRF instance deletes the OSPF instance if OSPF is disabled
- When you create a VRF instance, an OSPF instance is not automatically created. To activate OSPF on a VRF instance, first create an OSPF instance, and then enable OSPF
- You can configure a VRF so it can have IP interfaces with OSPF, RIP, static routes, and policies simultaneously
- Every IPv4 interface is a member of VRF 0 unless explicitly defined to belong to another VRF.

Virtualized Protocols

VRF Lite supports virtualization of the following IPv4 and IPv6 protocols and features. Use this table to find applicable VRF command and procedure information.

Table 24: Virtualized IPv4 Protocols and Documentation

Virtualized IPv4 protocol or feature	Where to find information
ARP	Configuring IPv4 Routing for VOSS
BGP	Configuring BGP Services for VOSS
Circuitless IP	Configuring IPv4 Routing for VOSS
DHCP	Configuring IPv4 Routing for VOSS
IGMP	Configuring IP Multicast Routing Protocols for VOSS
OSPF	Configuring OSPF and RIP for VOSS
RIP	Configuring OSPF and RIP for VOSS
Route policies	Configuring IPv4 Routing for VOSS
Route preferences	Configuring IPv4 Routing for VOSS
Router Discovery	Configuring IPv4 Routing for VOSS
Static routes	Configuring IPv4 Routing for VOSS
User Datagram Protocol (UDP)	Configuring IPv4 Routing for VOSS
VLAN	Configuring VLANs, Spanning Tree, and NLB for VOSS.
VRRP	Configuring IPv4 Routing for VOSS

Table 25: Virtualized IPv6 Protocols and Documentation

Virtualized IPv6 protocol or feature	Where to find information
BGP	Configuring BGP Services for VOSS
IPv6 Interfaces and IPv6 Static Routes	Configuring IPv6 Routing for VOSS

Virtualized IPv6 protocol or feature	Where to find information
ECMP and Alternative Route	Configuring IPv6 Routing for VOSS
OSPF	Configuring IPv6 Routing for VOSS
Route redistribution for static and direct routes	Configuring IPv6 Routing for VOSS
VRRPv3	Configuring IPv6 Routing for VOSS
DHCP Relay	Configuring IPv6 Routing for VOSS
IPv6 Reverse Path Forwarding	Configuring IPv6 Routing for VOSS
ICMP Ping & Traceroute	Configuring IPv6 Routing for VOSS
ISIS Accept Policies	Configuring Fabric Layer 3 Services for VOSS

VRF Lite Configuration using CLI

Use the procedures in this section to configure VRFs using CLI.

Creating a VRF Instance

About this task

Create a VRF instance to provide a virtual routing interface for a user.

For more information on route scaling, see Release Notes for VOSS.



Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a VRF instance and specify a VRF name:

```
ip vrf WORD<1-16>
```

- 3. Configure the maximum number of routes:
 - For IPv4:

```
ip vrf WORD<1-16> max-routes <0-15744 | 0-15488 | 0-256000>
```

• For IPv6:

```
ip vrf WORD<1-16> ipv6-max-routes <0-7744 | 0-7872>
```

- 4. Enable max-routes traps:
 - For IPv4:

```
ip vrf WORD<1-16> max-routes-trap enable
```

• For IPv6:

ip vrf WORD<1-16> ipv6-max-routes-trap enable



The maximum number of IPv6 routes for the Global Router are fixed and cannot be changed.

5. Enter VRF Router Configuration mode:

```
router vrf WORD<1-16>
```

6. Configure the IP routing protocol triggers for the VRF:

Use one of the following commands on your switch:

- ip bgp
 - ip bgp creates both ipv4 and ipv6 instances.
- ip ospf

Use ipv6 ospf to create an OSPFv3 instance.

• ip rip

RIPng is not virtualized, hence the IPv6 configuration is not applicable here.

Note:

You cannot configure BGP, OSPF, or RIP on a VRF instance unless you first configure the routing protocol trigger.

7. Ensure that the instance is configured correctly:

```
show ip vrf WORD<1-16>
```

Example

Create a VRF instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip vrf vrfRED
```

Configure the maximum number of IPv4 routes and enable max-routes traps.

```
Switch:1(config) #ip vrf vrfRED max-routes 12000
Switch:1(config) #ip vrf vrfRED max-routes-trap enable
```

Enter Router Configuration mode and configure the routing protocol triggers for the VRF:

```
Switch:1(config) #router vrf vrfRED
Switch:1(router-vrf) #ip bgp
Switch:1(router-vrf)#ip ospf
Switch:1(router-vrf) #ip rip
```

To Configure OSPFv3 instance for the VRF:

```
Switch:1(config) #router vrf vrfRED
Switch:1(router-vrf)#ipv6 ospf
```

Exit to Global configuration mode:

Switch:1(router-vrf)#exit

Configure the maximum number of IPv6 routes and enable IPv6 max-routes traps.

```
Switch:1(config) #ip vrf vrfRED ipv6-max-routes 7700
Switch:1(config) #ip vrf vrfRED ipv6-max-routes-trap enable
```

Ensure that the instance is configured correctly:

Switch:1#	Switch:1#show ip vrf vrfRED								
	VRF INFORMATION								
VRF COUNT	OSPF COUNT	RIP COUI	VT	BGP COUNT	PIM COUNT	ARI COU		PIM6 COUNT	OSPFv3 COUNT
1	1	1		1	1	7		1	1
VRF NAME	VRF ID	OSPF	RIP	BGP	PIM	VLAN COUNT	ARP COUNT	PIM6	OSPFv3
vrfRED	3	TRUE	TRUE	TRUE	TRUE	0	7	TRUE	FALSE
1 out of	l out of 1 Total Num of VRF Entries displayed.								

Variable definitions

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



Note:

Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Table 26: Variable definitions

Variable	Value
Depending on your hardware platform: max-routes <0-15488 0-15744	Configures the maximum number of IPv4 routes allowed for the VRF, which is 15488, 15744, or 252000, depending on your hardware platform.
0-252000>	The default value is 10000, except for the GlobalRouter, which is 15488, 15744, or 252000, depending on your hardware platform.
ipv6-max-routes <0-7744 0-7872>	Configures the maximum number of IPv6 routes allowed for the VRF, which is 7744 or 7872, depending on your hardware platform.

Variable	Value
	The default value is 5000.
max-routes-trap enable	Enables SNMP traps after the maximum number of IPv4 routes are reached.
ipv6-max-routes-trap enable	Enables SNMP trap generation based on the configured number of maximum IPv6 routes. The default is enabled.
name WORD<0-16>	Renames the VRF instance.
vrf-trap	Enables the device to send VRF-related traps.

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

Table 27: Variable definitions

Variable	Value
max-routes [vrfids WORD<0-512>]	Displays the maximum number of routes for the specified VRFs.
[WORD <1–16>]	vrfids WORD<0-512> specifies a list of VRFs by VRF IDs.
	WORD<1-16> specifies a VRF by name.
vrfids WORD<0-512>	Specifies a list of VRFs by VRF IDs.
WORD<1-16>	Specifies a VRF by name.

Associating a VLAN or port with a VRF instance

You can assign a VRF instance to a port or VLAN. You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IP address. You can configure the IP address after you associate the port and VRF instance.

Before you begin

• Ensure the VRF is already configured.

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

2. Associate the port or VLAN with a VRF instance:

```
vrf WORD<1-16>
```

Example

Switch:1> enable
Switch:1# configure terminal

Create a VRF named Two:

Switch:1(config-if) # ip vrf Two

Create a VLAN of type byport:

Switch:1(config-if) # vlan create 33 name vlan-30 type port-mstprstp 0

Enter VLAN Interface Configuration mode:

Switch:1(config-if) # interface vlan 33

Assign the VLAN to VRF Two:

Switch:1(config-if) # vrf Two

Give the VLAN an IP address:

Switch:1(config-if) # ip address 192.0.2.1 255.255.255.0

Enter VRF configuration mode:

Switch:1(config-if)# router vrf Two

Variable definitions

Use the data in the following table to use the ${\tt vrf}$ command.

Table 28: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies the VRF instance by name.

Creating an IPv6 VPN instance

Before you begin

The VRF must exist.

About this task

Create an IPv6 VPN instance to advertise IPv6 routes from a VRF to Shortest Path Bridging MAC (SPBM) network. For more information about Layer 3 Virtual Services Networks (VSNs) and SPBM, see Configuring Fabric Layer 3 Services for VOSS.

Procedure

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

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

2. Create an IPv6 VPN instance:

```
ipv6 ipvpn
```

3. Assign a service instance identifier (I-SID) to the IPv6 VPN:

```
i-sid < 0-16777215 >
```

4. Enable IPv6 VPN:

```
ipv6 ipvpn enable
```

5. Display all IPv6 VPNs:

```
show ipv6 ipvpn [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example

Create and enable IPv6 VPN instance:

```
Switch:1>enable
Switch: 1#configure terminal
Switch:1(config) #router vrf vrfred
Switch:1(router-vrf)#ipv6 ipvpn
Switch:1(router-vrf)#i-sid 100
Switch:1(router-vrf) #ipv6 ipvpn enable
Switch:1(router-vrf) #show ipv6 ipvpn
Switch:1(router-vrf) #show ipv6 ipvpn
   VRF Name
                       : vrfred
   Ipv6 Ipvpn-state : enabled
   Ipv4 Ipvpn-state : disabled
   I-sid
                         : 5555
Total active Ipv6 L3 VSN: 1
1 out of 3 Total Num of VRF Entries displayed.
```

Variable definitions

Use the data in the following table to configure the ipv6 ipvpn command.

Variable	Value
enable	Enables IPv6 IPVPN. The default is disabled.

Use the data in the following table to configure the i-sid command.

Variable	Value
<0–16777215>	Assigns an I-SID to the VRF being configured.

Use the data in the following table to configure the show ipv6 ipvpn command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

Enabling IPv6 trap notifications

About this task

Perform this procedure to enable SNMP traps when maximum number of IPv6 routes are reached.



Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable max-routes trap:

```
ip vrf WORD<1-16> ipv6-max-routes-trap enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #ip vrf vrfRED ipv6-max-routes-trap enable
```

Variable definitions

Use the data in the following table to use the ipv6-max-routes-trap command.

Variable	Value	
enable	Enables SNMP trap generation based on the configured number of	
	maximum IPv6 routes. The default is enabled.	

Displaying IPv6 max-routes information

About this task

Perform this procedure to display the maximum IPv6 routes configured.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display max-routes information:

```
show ip vrf ipv6-max-routes [vrfids WORD<0-512> | WORD<1-16>]
```

Example

```
Switch:1#show ip vrf ipv6-max-routes

VRF Specific Configuration

VRF-ID VRF-NAME CONTEXT-NAME IPV6-MAX-ROUTES IPV6-MAX-ROUTES-TRAP VRF-TRAP

0 GlobalRouter 8000 enable NA
1 vrfred vrf1 5000 enable enable
2 vrfblue vrf2 5000 enable enable
3 out of 3 Total Num of VRF Entries displayed.
```

VRF Lite Configuration using EDM

Use the procedures in this section to configure VRFs using EDM.

Configuring a VRF instance

About this task

Configure a VRF instance to provide a virtual routing interface for a user.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VRF**.
- 2. Click VRF.
- 3. Click the VRF tab.
- 4. Click Insert.
- 5. Specify the VRF ID.
- 6. Name the VRF instance.
- 7. Configure VRF Lite-related traps.
- 8. Configure the other parameters as required.
- 9. Click Insert.

VRF field descriptions

Use the data in the following table to help you use the VRF tab.

Name	Description	
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.	

Name	Description
Name	Names the VRF instance.
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB port management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is enabled.
MaxRoutes	Configures the maximum number of routes allowed for the VRF, which is 15488 or 15744, depending on your hardware platform.
	The default value is 10000, except for the GlobalRouter, which is 15488 or 15744, depending on your hardware platform.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. You can act upon multiple RPs simultaneously. You can also use this option to bring individual RPs up in steps.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is enabled.
Ipv6MaxRoutes	Configures the maximum number of IPv6 routes allowed for the VRF, which is 7744 or 7872, depending on your hardware platform.
	The default value is 5000.
Ipv6MaxRoutesTrapEnable	Enables SNMP trap generation after the maximum number of IPv6 routes are reached.
	The default is enabled.

Selecting and launching a VRF context view

About this task

Use this procedure to switch to another VRF context view when you use the embedded EDM. GlobalRouter is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open only five tabs for each EDM session.

Important:

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

Note:

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned. If you use a tenant VRF view, it is recommended to use the applicable EDM plugin with COM to access EDM. COM provides VRF mapping and Role-Based Access Control.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > VRF Context View**.
- 2. Click Set VRF Context View.
- 3. Click the VRF tab.
- 4. Select a context to view.
- 5. Click Launch VRF Context view.

A new browser tab opens containing the selected VRF view

VRF field descriptions

Use the descriptions in the following table to use the **VRF** tab.

Name	Description
Id	Shows the unique VRF ID.
Name	Shows the name of the virtual router.
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB port management.

Creating an IPv6 VPN instance on a VRF

Create an IPv6 VPN instance to advertise IPv6 routes from a VRF to Shortest Path Bridging MAC (SPBM) network.

For more information about Layer 3 Virtual Services Networks (VSNs) and SPBM, see <u>Configuring Fabric Layer 3 Services for VOSS</u>.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- · The VRF must exist.

Procedure

- 1. In the navigation tree, expand the following folders: Configuration > IPv6.
- 2. Click IPv6-VPN.
- 3. Click the VPN tab.
- 4. Click Insert.
- 5. Click the ellipsis [...], and select a VRF from the list.
- 6. Click OK.
- 7. Click Insert.

- 8. In the **IsidNumber** column, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IPv6-VPN.
- 9. In the **Enable** column, double-click the value and select **true** or **false** from the drop-down list.
- 10. Click Apply.

VPN field descriptions

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

Name	Description
Vrfld	Specifies the ID of the VRF to configure.
Enable	Enables or disables the IPv6 VPN instance on the VRF. The default is disabled.
IsidNumber	Specifies the I-SID to associate with the IPv6 VPN. By default, no I-SID is assigned.

Associating a port to a VRF instance

About this task

Associate a port to a Virtual Router Forwarding (VRF) instance so that the port becomes a member of the VRF instance.

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the GlobalRouter, VRF 0, by default.

Procedure

- 1. In the **Device Physical** View tab, select a port.
- 2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
- 3. Click General.
- 4. Click the **VRF** tab.
- 5. To the right of the **BrouterVrfld** box, click the ellipsis (...) button.
- 6. In the BrouterVrfld dialog box, select the required VRF.
- 7. Click OK.
- 8. Click Apply.

VRF field descriptions

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

Name	Description
Vrflds	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the port is associated.
BrouterVrfld	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Associating an Insight Port to a VRF Instance

About this task

Perform this procedure to associate an Insight port to a Virtual Router Forwarding (VRF) instance.



You can associate a VRF instance to an Insight port after you configure the VRF. By default, the Insight ports are associated to the GlobalRouter.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit > Insight Port**.
- 2. Click the Insight port you want to configure.
- 3. Click the VRF tab.
- 4. **(Optional)** In the **VrfNames** field, Click the **ShowAll** button to view the VRF instances the Insight Port is associated with.
- 5. In the **BrouterVrfld** field, click the ellipsis (...) button, and select the required VRF instance(s).
- 6. Click Ok.
- 7. Click Apply.

VRF Field Descriptions

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

Name	Description
Vrflds	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the Insight port is associated with.
BrouterVrfld	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Appendix A: ICMPv6 type and code

The Internet Control Message Protocol (ICMPv6) uses many messages identified by a type and code field (see RFC 4443). Error messages use message types 0 to 127. Informational messages use message types 128 to 255. The following table provides the type and code reference.

Table 29: ICMPv6 type and code details

Туре	Name	Code	Reference
1	Destination Unreachable	0—no route to destination	RFC 4443
		1—communication with destination administratively prohibited	
		2—(not assigned)	
		3—address unreachable	
		4—port unreachable	
2	Packet Too Big	N/A	RFC 4443
3	Time Exceeded	0—hop limit exceeded in transit	RFC 4443
		1—fragment reassembly time exceeded	
4	Parameter Problem	0—erroneous header field encountered	RFC 4443
		1—unrecognized Next Header type encountered	
		2—unrecognized IPv6 option encountered	
128	Echo Request	N/A	RFC 4443
129	Echo Reply	N/A	RFC 4443
130	Multicast Listener Query	N/A	
131	Multicast Listener Report	N/A	
132	Multicast Listener Done	N/A	
133	Router Solicitation	N/A	RFC 4861

Туре	Name	Code	Reference
134	Router Advertisement	N/A	RFC 4861
135	Neighbor Solicitation	N/A	RFC 4861
136	Neighbor Advertisement	N/A	RFC 4861
137	Redirect Message	N/A	RFC 4861
138	Router Renumbering	0—router renumbering command 1—router renumbering result 255—sequence number reset	
139	ICMP Node Information Query	N/A	
140	ICMP Node Information Response	N/A	
141	Inverse neighbor discovery Solicitation Message	N/A	RFC 3122
142	Inverse neighbor discovery Advertisement Message	N/A	RFC 3122
143	Version 2 Multicast Listener Report	N/A	RFC 3810
144	Home Agent Address Discovery Request Message	N/A	RFC 3775
145	Home Agent Address Discovery Reply Message	N/A	RFC 3775
146	Mobile Prefix Solicitation	N/A	RFC 3775
147	Mobile Prefix Advertisement	N/A	RFC 3775

Glossary

All_DHCP_Relay_Ag ents_and_Servers

(FF02::1:2)

A link-scoped multicast address used by a client to communicate with neighboring relay agents and servers. All servers and relay agents are members of this multicast group.

Dual stack

Supports both the IPv4 and IPv6 protocol.

A collection of ports in one spanning-tree instance.

Extended Unique Identifier (EUI)

A 64-bit format used in assigning addresses automatically to IPv6

interfaces.

route table manager (RTM)

(13110)

Determines the best route to a destination based on reachability, route

preference, and cost.

Spanning Tree Group (STG)

(316)

stateless address autoconfiguration

(SLAAC)

Uses a mathematical equation to automatically configure and assign IPv6 addresses to hosts or nodes on a network. RFC 4862 describes SLAAC.