

Configuring IPv6 Routing on Avaya Virtual Services Platform 9000

Release 4.1 NN46250-509 Issue 06.01 October 2015

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: <u>http://support.avaya.com</u> or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE. HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC. ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

Licence types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

"Heritage Nortel Software" means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at <u>http://support.avaya.com/ LicenseInfo</u> under the link "Heritage Nortel Products" or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya's website at: <u>http://</u> <u>support.avaya.com/Copyright</u> or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: <u>http://support.avaya.com</u>, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux $^{\ensuremath{\mathbb{R}}}$ is the registered trademark of Linus Torvalds in the U.S. and other countries.

Contents

Chapter 1: Introduction	8
Purpose	8
Related resources	8
Documentation	8
Training	8
Viewing Avaya Mentor videos	9
Support	9
Searching a documentation collection	10
Chapter 2: New in this release	11
Features	11
Other changes	12
Chapter 3: IPv6 basics	13
Origins of IPv6	
Advantages of IPv6	
Comparison of IPv4 and IPv6	
IPv6 packet	15
IPv6 header	15
IPv6 extension headers	16
IPv6 address component summary	18
IPv6 address formats	19
Address types	19
IP address prefix	24
Interface ID	24
How to write an IPv6 address	
IPv4 compatible addresses	25
ICMPv6	
Path MTU discovery	
Routing	
IPv6 basic configuration using ACLI	
Configuring an IPv6 static neighbor address	
Configuring an IPv6 interface	
Assigning IPv6 addresses to a brouter port or VLAN	
Viewing global IPv6 information	
Creating IPv6 static routes	
Viewing routes information	
IPv6 basic configuration using EDM	
Configuring IPv6 forwarding	
Configuring an IPv6 interface	
Assigning IPv6 addresses to interfaces	51

Creating IPv6 static routes	. 53
Viewing route information	. 54
Chapter 4: Neighbor discovery	. 56
Neighbor discovery	. 56
Host autoconfiguration	. 61
Neighbor discovery configuration using ACLI	. 62
Configuring an IPv6 discovery prefix	. 62
Configuring route advertisement	. 64
Configuring the neighbor cache	. 69
Viewing cached destination information	. 71
Neighbor discovery configuration using EDM	. 72
Configuring an IPv6 discovery prefix	. 72
Configuring route advertisement	. 75
Configuring the neighbor cache	. 77
Viewing cached destination information	
Chapter 5: DHCP Relay	. 80
DHCP Relay	
DHCP Relay configuration using ACLI	
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	
Chapter 6: OSPFv3	
OSPFv3	
OSPFv3 configuration using ACLI	
Configuring OSPF globally	
Creating an OSPF area	
Creating OSPF area ranges	
Creating an OSPF virtual link	
Configuring OSPF on a port or VLAN	. 99
Viewing OSPFv3 information	102
Adding an NBMA neighbor	105
Enabling OSPF route redistribution	
Viewing the status of OSPFv3 redistribution	
OSPFv3 configuration using EDM	108
Configuring OSPF globally	
Creating an OSPF area	
Creating OSPF area ranges	111
Creating an OSPF virtual link	112

Creating an OSPF interface	114		
Creating an OSPF VLAN interface			
•			
•			
•			
Viewing OSPF neighbors	126		
Viewing virtual neighbors	127		
•			
Adding an NBMA neighbor. 11 Enabling OSPF route redistribution. 11 Modifying an OSPF interface. 12 Viewing OSPF neighbors. 12 Viewing virtual neighbors. 12 Viewing virtual neighbors. 12 Viewing virtual neighbors. 12 Chapter 7: VRRP. 11 VRRP. 11 VRRP. 12 VRRP configuration using ACLI. 12 Configuring the VRRP interface. 11 Configuring VRRP notification control. 12 Configuring VRRP notification control. 13 Configuring VRRP notification control. 14 Configuring VRRP for an interface. 14 Configuring VRRP for a VLAN. 14 Configuring VRRP notification control. 14 Configuring RSMLT 14 Configuring RSMLT 14 Configuring RSMLT on a VLAN. 14 Configuring RSMLT on a VLAN. 14 <			
0 0			
Viewing the AS-scope Link-state database. 1 Viewing the link-scope LSDB. 1 Adding an NBMA neighbor. 1 Enabling OSPF route redistribution. 1 Modifying an OSPF interface. 1 Viewing virtual neighbors. 1 Viewing virtual neighbors. 1 Viewing virtual neighbors. 1 Viewing virtual neighbors. 1 VRP. 1 VRRP configuration using ACLI. 1 Configuring the VRRP interface. 1 Viewing VRRP information. 1 Configuring VRRP notification control. 1 Configuring VRRP for an interface. 1 VRRP configuration using EDM. 1 Configuring VRRP for a ninterface. 1 Configuring VRRP for a VLAN. 1 Configuring VRRP for a VLAN. 1 Configuring RSMLT Configuration using ACLI. 1 Configuring RSMLT Information. 1 Configuring RSMLT export. 1 Configuring RSMLT on a VLAN. 1 Configuring RSMLT on a VLAN. 1 Enabling RSMLT fon a VLAN. 1			
• •			
Configuring additional VRRP parameters for an interface.138VRRP configuration using EDM.140Configuring VRRP for an interface.140Configuring VRRP for a VLAN.143Configuring VRRP notification control.145Configuring additional addresses on the VRRP interface.146Chapter 8: RSMLT.148RSMLT.148RSMLT configuration using ACLI.149Configuring RSMLT on a VLAN.149Enabling RSMLT Edge support.151			
• • •			
5			
· ·			
Chapter 7: VRRP. 12 VRRP. 12 VRRP configuration using ACLI. 13 Configuring the VRRP interface. 13 Viewing VRRP information. 13 Configuring VRRP notification control. 13 Configuring additional VRRP parameters for an interface. 13 VRRP configuration using EDM. 14 Configuring VRRP for an interface. 14 Configuring VRRP for a vLAN. 14 Configuring VRRP notification control. 14 Configuring VRRP notification control. 14 Configuring VRRP for a vLAN. 14 Configuring VRRP notification control. 14 Configuring RRP notification control. 14 Configuring RSMLT 14 RSMLT 14 RSMLT 14 RSMLT configuration using ACLI. 14 Configuring RSMLT on a VLAN. 15 Viewing RSMLT information. 15 Configuring RSMLT on a VLAN. 15			
•			
•			
•			
•			
enapter ret in to configuration examples	100		

OSPFv3	166
Dual stack core	168
Appendix A: ICMPv6 type and code	177
Glossary	

Chapter 1: Introduction

Purpose

This document provides conceptual and procedural information to configure IPv6 routing operations on the Avaya Virtual Services Platform 9000. 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), and Routed Split MultiLink Trunking (RSMLT).

For information about BGP for IPv6, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

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

For detailed information about SNMPv3, see *Configuring Security on Avaya Virtual Services Platform 9000,* NN46250-601.

For detailed information about SSH, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

For information on prefix lists for IPv4 and IPv6, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Related resources

Documentation

See Documentation Reference for Avaya Virtual Services Platform 9000, NN46250-100 for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at <u>http://avaya-learning.com/</u>.

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to <u>http://support.avaya.com</u> and perform one of the following actions:
 - In Search, type Avaya Mentor Videos to see a list of the available videos.
 - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to <u>www.youtube.com/AvayaMentor</u> and perform one of the following actions:
 - Enter a key word or key words in the **Search Channel** to search for a specific product or topic.
 - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

😵 Note:

Videos are not available for all products.

Support

Go to the Avaya Support website at <u>http://support.avaya.com</u> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Searching a documentation collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

- 1. Extract the document collection zip file into a folder.
- 2. Navigate to the folder that contains the extracted files and open the file named cproduct_name_release>.pdx.
- 3. In the Search dialog box, select the option **In the index named** cproduct_name_release>.pdx.
- 4. Enter a search word or phrase.
- 5. Select any of the following to narrow your search:
 - Whole Words Only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
- 6. Click Search.

The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Chapter 2: New in this release

The following details what is new in *Configuring IPv6 Routing on Avaya Virtual Services Platform 9000,* NN46250-509, for Release 4.1.

Features

See the following sections for information about feature changes.

IPv6 support

Release 4.1 supports IPv6 configuration using Avaya Command Line Interface (ACLI) and Enterprise Device Manager (EDM). Border Gateway Protocol Plus (BGP+), IPv6 tunnels, IPv6 Shortcuts, and IPv6 filters are not supported in this release.

Licensing

Release 4.1 introduces the Product Licensing and Delivery System (PLDS) as the license order, delivery, and management tool. Release 4.1 includes features in either the Base License or the Premier License. Features that were included in the Advanced License, in previous releases, are now included in the Base License.

The following IPv6 features are now included in the Base License:

- Neighbor Discovery
- DHCP Relay
- OSPFv3
- VRRP
- RSMLT

For more information, see:

- IPv6 basics on page 13.
- Neighbor discovery on page 56.
- DHCP Relay on page 80.
- <u>OSPFv3</u> on page 90.
- <u>VRRP</u> on page 129.
- <u>RSMLT</u> on page 148.
- Viewing IPv6 connections on page 159.

For more information on PLDS and licensing conceptual and configuration information, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Neighbors tab

Release 4.1 updates the **Neighbors** tab to remove the **Bmac** field. For more information, see <u>Configuring the neighbor cache</u> on page 77.

RSMLT holddown timer

Release 4.1 updates RSMLT holddown timer information. You cannot clear the static IPv6 routes during the RSMLT holddown timer. The RSMLT holddown timer defines how long the recovering or restarting system remains in non-Layer 3 forwarding mode for the peer route MAC address.

For more information, see:

- Configuring RSMLT on a VLAN on page 149.
- RSMLT field descriptions on page 154.

Other changes

There are no other changes in this release.

Chapter 3: IPv6 basics

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

Before you begin

Avaya includes the IPv6 features in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

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.

Note:

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

Table 1: IPv4 and IPv6 key differences compared

Feature	IPv4	IPv6
Address length	32 bits	128 bits
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

Table continues...

Feature IPv4		IPv6
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

😵 Note:

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 2: 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)	8

Table continues...

Field	Size in bits
Source address	128
Destination address	128

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

	Ver	DS byte	Flow Label		
	Pay	/load Leng	ngth Next header Hop Limit		
40 Byte	Source Address Destination Address				

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.

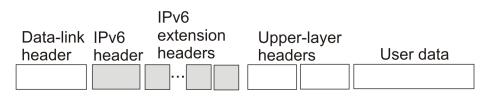


Figure 3: IPv6 header with extension headers

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 3: 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
- Multicast
- Anycast

Unicast:

Unicast addresses provide one-to-one communication.

Unicast addresses can be

- Global
- · Link local
- Special, for example: Unspecified and Loopback

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.

The switch supports Subnet-Router Anycast: example <prefix>::0.

Broadcast:

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

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 19

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

Unicast addresses

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.

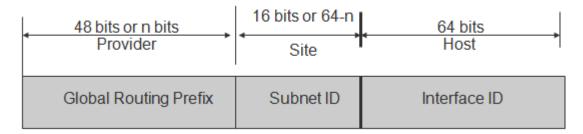


Figure 5: Uicast global address parts

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.

← 10 bits ← FE80	s4 bits ↓	← 64 bits
1111 1110 10	000000	Interface ID

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).

Equivalent to the IPv4 unspecified address 0.0.0.0, represented as 0:0:0:0:0:0:0:0:0:0:0::0: or ::1; an IPv6 host that does not have a valid address uses the unspecified address as its source address when it sends a packet to discover whether an address is used by another node (during the boot process when the host requests address configuration information).

😵 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.

IPv6 Unicast address example:

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

IPv6 Link-local Unicast address example:

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

Multicast addresses

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
1111111	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 4: Well known multicast IPv6 addresses

Name	Address
All Nodes	FF02:0:0:0:0:0:1
All Routers	FF02:0:0:0:0:0:0:2
OSPFIGP	FF02:0:0:0:0:0:0:5
OSPFIGP Designated Routers	FF02:0:0:0:0:0:06
All PIM Routers	FF02:0:0:0:0:0:D
VRRP	FF02:0:0:0:0:0:0:12
All MLDv2–capable routers	FF02:0:0:0:0:0:0:16
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.

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.

Difference between multicast and anycast

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

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 61.

How to write an IPv6 address

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

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:XXXX

Hexadecimal IPv6 address representations:

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, where n is the hexadecimal representation of 16 bits in the address; for example, FF01:0:0:0:0:0:0:0:43.

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—FF01::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:

```
2001:10F2:0000:0000:25AB:0000:0000:0001
2001:10F2:0:0:25AB:0:0:1
2001:10F2:0:0:25AB::1
2001:10F2::25AB:0:0:1
```

IPv4 compatible addresses

IPv6 nodes that need to operate with IPv4 nodes use the IPv4-compatible address, which includes an IPv4 address in the low-order 32 bits.

The following figure shows the format of an IPv4-compatible address.

96 bits	32 bits
0000 0000 0000 0000 0000 0000	IPv4 address

Figure 8: IPv4–compatible address format

An IPv4-compatible address combines hexadecimal and decimal values as x:x:x:x:x:d.d.d.d.

In the previous address format x:x:x:x:x is a hexadecimal representation of the 6 high-order 16-bit pieces of the address, and d.d.d is a decimal representation of the four 8-bit pieces of the address.

Examples of an IPv4 compatible address:

0:0:0:0:0:0:13.1.68.3

::13.1.68.3.

The following figure illustrates the IPv4 device address 222.1.41.90 converted to an IPv4-mapped IPv6 address; note the IPv6 address mixed notation and compressed IPv6 address formats.

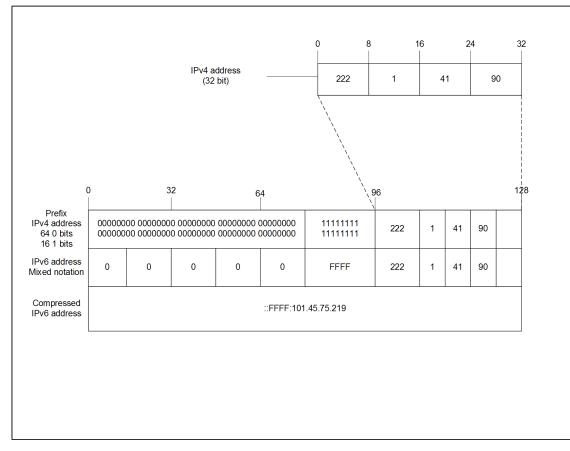


Figure 9: IPv4–mapped IPv6 address

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

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.

Routing

A routing table is present on all nodes.

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

😵 Note:

The switch requires routing protocols, such as OSPFv3 or BGP, 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. Offlink 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 OSPFv3 and BGP as the IPv6 routing protocols.

You can redistribute directly connected routes, IPv6 static routes, and IPv6 BGP routes into OSPFv3.

You can redistribute directly connected routes, IPv6 static routes, and IPv6 OSPFv3 routes into BGP.

This document focuses on OSPFv3. For information about OSPFv2, see *Configuring OSPF and RIP* on Avaya Virtual Services Platform 9000, NN46250-506.

For information about BGP for IPv6, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507

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.

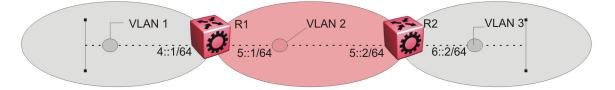


Figure 10: IPv6 routing between VLANs

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 **TryToResolve** 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 VLAN, or port. If you use outgoing VLAN, or port, the implied nexthop value is 0::0.

When you configure static routes with a link-local nexthop, you must also specify the outgoing 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 5: 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
	You can use one ACLI command to adminstratively enable or disable all static routes as follows ipv6
	route static enable.

Table continues...

Action	Result
	You can administratively disable all routes but preserve the static route configuration when you use the following ACLI 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.

The local-nexthop flag is not required for Pv6.

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

The maximum number of entries is 10,000. 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.

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.

IPv6 basic configuration using ACLI

Configuring an IPv6 static neighbor address

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

Before you begin

• Enable IPv6 forwarding.

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

Because IPv6 forwarding is disabled by default, you can only use local IPv6 connections and traffic does not traverse an IPv6 network.

To configure an IPv6 static neighbor address you must enable IPv6 forwarding.

Procedure

1. Enter Global Configuration mode:

enable

configure terminal

2. Enable forwarding:

ipv6 forwarding

3. Configure an IPv6 neighbor address:

```
ipv6 neighbor WORD<0-128> port {slot/port} mac
<0x00:0x00:0x00:0x00:0x00:0x00> [vlan <1-4084>
```

- 4. 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 the ICMP error interval:

ipv6 icmp error-interval <0-2147483647>

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

c. Configure the ICMP error quota:

ipv6 icmp error-quota <0-2000000>

The default is 50.

d. Enable ICMP redirect messages:

ipv6 icmp redirect-msg

The default is disabled.

e. Enable ICMP network unreachable messages:

ipv6 icmp unreach-msg

The default is disabled.

Example

Enable IPv6 forwarding:

Switch:1(config)#ipv6 forwarding

Add an IPv6 neighbor for a brouter port:

```
Switch:1(config)#ipv6 neighbor 3000:0:0:0:0:0:0:2 port 10/11 mac 00:0c: 42:07:35:90
```

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 disabled.
hop-limit <0-255>	Configures the hop limit. The default is 64.
icmp error interval <0-2147483647>	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000. An entry of 0 seconds results in no sent ICMPv6 error messages
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 redirect-msg	Enables ICMP redirect messages. The default is disabled.
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 <1-4084>]	Creates a static IPv6 neighbor with the following variables:
	 WORD<0-128> specifies the IPv6 addressof the neighbor in hexadecimal colon format.
	 {slot/port} specifies the brouter port to use for the neighbor. Identifies a single slot and port.
	 mac <0x00:0x00:0x00:0x00:0x00:0x00> specifies the MAC address of the neighbor.
	 vlan <1-4084> specifies the VLAN ID to use for the neighbor.
	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
	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, protocol-based, and MAC-source-based VLANs.

For information about how to configure VLANs, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500 and *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000,* NN46250-503.

Procedure

1. Enter Interface Configuration mode:

enable

configure terminal

```
interface GigabitEthernet {slot/port[-slot/port][,...]} OF interface
vlan <1-4084>
```

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 maximum transmission unit (MTU):

ipv6 interface mtu <1280-9500>

e. Configure an interface description:

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

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

ipv6 interface reachable-time <1-3600000>

g. 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>

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

ipv6 interface vlan <1-4084>

- i. 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** ACLI 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.
mtu <1280–9500>	Configures the maximum transmission unit for the interface: 1280–1500, 1850, or 9500. This value

Table continues...

Variable	Value
	must be the same for all addresses defined on this interface.
	The default is 1500.
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.
	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 <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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 <exist-only strict></exist-only strict>	Specifies the mode for reverse path checking.
	In exist-only mode, reverse path checking checks whether the source address for the incoming packet exists in the routing table. If the source entry is found, the packet is forwarded as usual; otherwise, the packet is discarded.
	In strict mode, reverse path checking checks whether the source address for the incoming packet exists in the routing table. If the source entry is found, reverse path checking further checks if the source interface address matches the packet incoming interface address. If they match, the packet is forwarded as usual, otherwise, the packet is discarded.

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

Variable	Value
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

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[-slot/port][,...]} OF interface
vlan <1-4084>
```

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:ffff/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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

Viewing global IPv6 information

Perform this procedure to view and manage general IPv6 information.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display IPv6 information for an interface:

```
show ipv6 interface [gigabitethernet {slot/port[-slot/port][,...]}]
[mgmtEthernet {slot/port[-slot/port][vlan <1-4084>]
```

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

show ipv6 address interface ip WORD<0-46>

4. Display IPv6 address information for the specified VLAN:

show ipv6 address interface vlan <1-4084>

5. Display the current state of IPv6 forwarding:

show ipv6 forwarding

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

show ipv6 global

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

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

8. Display IPv6 VLAN router advertisement information:

show ipv6 nd interface vlan [<1-4084>]

9. Display detailed information in IPv6 router advertisements:

show ipv6 nd-prefix detail

10. Display GbE interface information in IPv6 router advertisements:

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

11. Display VLAN interface information in IPv6 router advertisements:

show ipv6 nd-prefix interface vlan [<1-4084>]

12. Display VLAN information in IPv6 router advertisements:

show ipv6 nd-prefix vlan <1-4084>

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

show ipv6 neighbor interface gigabitethernet <slot/port>

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

show ipv6 neighbor type <1-4>

15. Display IPv6 neighbor information:

show ipv6 neighbor [WORD<0-46>]

Example

Switch:1>enable Switch:1#show i		e gigabit	Ethernet							
		t Ipv6 In								
IFINDX BROUTER INDX	PHYSICAL	ADMIN OPE STATE STA	R TYPE MTU	HOP		LE RETRANSM TIME			ING RPC H	
344 5/25 80:17	:7d:76:8a:12	enable	up ETHER 150	0 64	30000	1000	disable	enable	disable	existonly
			Ipv6 Address					-		
IPV6 ADDRESS			BROUTER		TYPE	ORIGIN	STATUS	-		
2011:beef:4004: fe80:0:0:0:8217		al2	5/25 5/25		UNICAST UNICAST	MANUAL LINKLAYER	PREFERREI PREFERREI			

1 out of 11 Total Num of Interface Entries displayed. 2 out of 22 Total Num of Address Entries displayed. Switch:1#show ipv6 interface mgmtEthernet Port Ipv6 Interface -----------IFINDX BROUTER PHYSICAL ADMIN OPER TYPE MTU HOP REACHABLE RETRANSMIT MCAST FORWARDING RPC RPCMODE INDX ADDRESS STATE STATE LMT TIME TIME STATUS ADMIN STATUS 64 1/1 80:17:7d:76:83:fd enable up ETHER 1500 64 30000 1000 disable N/A N/A N/A _____ Port Ipv6 Address BROUTER TYPE ORIGIN STATUS IPV6 ADDRESS 2013:47:17:120:2:0:0:108 1/1 UNICAST MANUAL PREFERRED fe80:0:0:0:8217:7dff:fe76:83fd 1/1 UNICAST LINKLAYER PREFERRED 1 out of 10 Total Num of Interface Entries displayed. 2 out of 20 Total Num of Address Entries displayed. Switch:1#show ipv6 interface vlan Vlan Ipv6 Interface IFINDX VLAN PHYSICAL ADMIN OPER TYPE MTU HOP REACHABLE RETRANSMIT MCAST FORWARDING RPC RPCMODE INDX ADDRESS STATE STATE LMT TIME TIME STATUS ADMIN STATUS 214810080:17:7d:76:8a:01enable upETHER150064300001000disable enabledisable existonly215811080:17:7d:76:8a:02enable upETHER150064300001000disable enabledisable existonly224820080:17:7d:76:8a:02enable upETHER150064300001000disable enabledisable existonly225821080:17:7d:76:8a:02enable upETHER150064300001000disable enabledisable existonly254850080:17:7d:76:8a:09enable upETHER150064300001000disable enabledisable existonly264860080:17:7d:76:8a:0aenable upETHER150064300001000disable enabledisable existonly294890080:17:7d:76:8a:0eenable upETHER150064300001000disable enabledisable existonly Vlan Ipv6 Address ------TPV6 ADDRESS VLAN-ID TYPE ORIGIN STATUS VLAN-ID TYPE ORIGIN STATUS 2011:beef:100:0:0:0:0:67 V-100 UNICAST MANUAL PREFERRED 2011:beef:110:0:0:0:0:67 V-100 UNICAST LINKLAYER PREFERRED 2011:beef:110:0:0:0:0:67 V-110 UNICAST MANUAL PREFERRED --More-- (q = quit) Switch:1#show ipv6 address interface gigabitethernet 5/25 Address Information ------_____ IPV6 ADDRESS VID/BID/TID TYPE ORIGIN STATUS -----2011:beef:4004:0:0:0:67 5/25 UNICAST MANUAL PREFERRED fe80:0:0:0:8217:7dff:fe76:8a12 5/25 UNICAST LINKLAYER PREFERRED 2 out of 22 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 LINKLAYER PREFERRED 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 Global forwarding : enable Switch:1#show ipv6 global forwarding forwarding default-hop-cnt number-of-interfaces icmp-error-interval icmp-error-quota icmp-unreach-msg : enable : 64 : 11 : 1000 : 50 : disable

<pre>intpr-cdirect-mag : diable static-pote-administatus : shall Suitch: Heiny typ6 ind interface gigblicthernet</pre>														
Port Ipv6 Md First Prove Model The Marked OTHER DAD-NE NTU HOF REACHABLE RETRANSHIT FLAG OTHER DAD-NE NTU HOF REACHABLE RETRANSHIT FLAG OTHER DAD-NE NTU HOF REACHABLE RETRANSHIT FLAG OTHER DAD-NE NTU HOF REACHABLE RETRANSHIT I out of 11 Total Num of Ipv6 ND Entires displayed. Note: (d) = Default, (d) = inherit. 1 out of 11 Total Num of Ipv6 ND Entires displayed. Vian Ipv6 Nd Vian Ipv6 Nd Total Num of Ipv6 ND Entires displayed. Vian Ipv6 Nd Vian Ipv6 Nd Colspan="2">Total Num of Ipv6 Nd Vian Ipv6 Nd Vian NUM-ON MAX-INT MIN-INFL IFETIME MANG OTHER DAD-NE NTU HOF MEACHABLE RETRANSMIT Inflorme 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2158 V-100 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2258 V-230 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2258 V-230 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 200 1800 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d) 2268 V-240 True 600 201 100 False False 1 0(d) 64(d) 0(d) 0(d		icm _i stat	p-redire	ect-msg te-admin-	status	: dis : ena	able ble							
Port Type Nd IPTD BER RR-ADV MAX-INF MIN-INF LIFETIME MANOTHER DAD-50 MUT DD PARAMENE REPRESENT DIAG 5723 True 600 200 1800 Paise Faise 1 0(d) 64(d) 0(d) 0(d) Nete: (a) = Set. (d) = Default. (i) = inhort. Vian Type Nd 1 out of 11 Total Num of Ipvs ND Entries displayed. Suitch:148.00 MIX-INF MIN-INF LIFETIME MANOTHER DAD-76 MUT MORE THAN THE TIME Suitch:148.00 Ipv6 ind interface vian Vian Type Nd Nd	Switch	n:1#sho	ow ipv6	nd inter	face gig	abitether	net							
IPTO BRE RTR-ADV MAX-INF MIN-INF LIFETIME MANAG OWER DAD-NE MUT HOP REACHABLE RETRANSHT INAG INNE INF										=====				
FLAG CONF LIMIT TIME TIME 344 5/23 True 600 Default, (i) heats false False 10(d) 64(d) 0(d) 0(d) Note: (a) Seg. (d) Default, (i) heats false False 10(d) 64(d) 0(d) 0(d) Switch:14show ipv6 nd interface vian														
Note: (a) = Set, (d) = Default, (i) = inherit. 1 out of 11 Total Num of 1906 ND Entries displayed. Switch1#show ipv6 nd interface vlan							FLAG		CONF	I	JIMIT TI	ME	TIME	
<pre>1 out of 11 Total Num of Ipv6 ND Entries displayed. Switch:14show ipv6 nd interface vlan</pre>	344 5	5/25	True	600	200	1800	False	False	1 0(d)	6	54(d) 0((d)	0(d)	
Vian Ipvé Nd														

2948 2011:beef:900:0:0:0:0/64 900 2592000 604800 1

7 out of 9 Total Num of Ipv6 ND prefix Entries displayed.

Switch:1#show ipv6 nd-prefix vlan 100

	Nd-Prefix Address Informati	on			
=====			=========	========	
INTF	IPV6	VLAN	VALID	PREF	EUI
INDEX	ADDRESS/PREFIX	ID	LIFE	LIFE	
2148	2011:beef:100:0:0:0:0:0/64	100	2592000	604800	1

Legend: EUI: eui-not-used(1), eui-used-with-ul-complement(2)eui-used-without-ul-complement(3)

Switch:1#show ipv6 neighbor interface gigabitethernet 5/25

Ne	eighbor	Informa	tion		
NET ADDRESS/ PHYSICAL ADDRESS	IPV6 INTF	PHYS INTF	TYPE	STATE	LAST UPD
2011:beef:4004:0:0:0:0:67/ fe80:0:0:0:8217:7dff:fe76:8a12/	5/25 5/25	5/25 5/25			4877 80:17:7d:76:8 4877 80:17:7d:76:8

2 out of 73 Total Num of Neighbor Entries displayed.

Switch:1#show ipv6 neighbor type 2

Neighbor Information

Neighbor Information						
NET ADDRESS/ PHYSICAL ADDRESS	IPV6 INTF	PHYS INTF	TYPE	STATE	LAST UPD	
2013:47:17:120:0:0:0:1/ 2013:47:17:120:0:0:0:2/ 2013:47:17:120:1:0:0:7/ 2013:47:17:120:1:0:0:23/ 2013:47:17:120:1:0:0:231/ 2013:47:17:120:1:0:0:233/ 2013:47:17:120:1:0:0:239/ 2013:47:17:120:1:0:0:243/	1/1 1/1 1/1 1/1 1/1 1/1 1/1 1/1	1/1 1/1 1/1 1/1 1/1 1/1 1/1 1/1 1/1	DYNAMIC DYNAMIC DYNAMIC DYNAMIC DYNAMIC DYNAMIC DYNAMIC DYNAMIC	STALE STALE STALE STALE STALE	5640 00:1d:af:64:a2:01 5170 00:18:b0:5a:92:01 5321 80:17:7d:76:63:fd 5126 00:24:7f:a1:63:fd 5398 80:17:7d:76:63:ff 5195 80:17:7d:75:93:ff 5207 80:17:7d:75:93:fd 5190 00:24:7f:a1:63:ff	

--More-- (q = quit)

Switch:1(config)#show ipv6 neighbor

	Neighbor	Informat	cion			
NET ADDRESS/ PHYSICAL ADDRESS		IPV6 INTF	PHYS INTF	TYPE	STATE	LAST UPD
2013:47:17:120:0:0:0:1/ 00:1d:af:64:a2:01		1/1	1/1	DYNAMIC	STALE	5681
2013:47:17:120:0:0:0:2/ 00:18:b0:5a:92:01		1/1	1/1	DYNAMIC	STALE	5170
2013:47:17:120:1:0:0:7/ 80:17:7d:76:63:fd		1/1	1/1	DYNAMIC	STALE	5321
2013:47:17:120:1:0:0:23/ 00:24:7f:a1:63:fd		1/1	1/1	DYNAMIC	STALE	5126
2013:47:17:120:1:0:0:231/ 80:17:7d:76:63:ff		1/1	1/1	DYNAMIC	STALE	5398
2013:47:17:120:1:0:0:233/ 80:17:7d:75:93:ff		1/1	1/1	DYNAMIC	STALE	5195
2013:47:17:120:1:0:0:239/ 80:17:7d:75:93:fd		1/1	1/1	DYNAMIC	STALE	5207
2013:47:17:120:1:0:0:243/ 00:24:7f:al:63:ff		1/1	1/1	DYNAMIC	STALE	5190

--More-- (q = quit)

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.

Variable	Value
neighbor [WORD<0-46>]	Specifies the IPv6 address of the neighbor.
{slot/port[-slot/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).
type <1-4>	Specifies the neighbor type as one of the following:
	• 1-other
	• 2-dynamic
	• 3-static
	• 4-local
vlan <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Creating IPv6 static routes

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

Before you begin

• Enable IPv6 forwarding.

Procedure

1. Enter Global Configuration mode:

enable

configure terminal

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>] [{slot/port} ] [vlan <1-4084>]
```

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

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.
next-hop <i>Word <0–46></i>	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}	Identifies a single slot and port.
vlan <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Viewing routes information

View routes information to view the current configuration.

Procedure

1. Enter Privileged EXEC mode:

enable

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

show ipv6 route count-summary

- 3. Show route information for a destination: show ipv6 route dest WORD<0-46>
- 4. Show route information for a port:

show ipv6 route gigabitethernet {slot/port}

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

show ipv6 route next-hop WORD<0-46>

6. Show route information for a static route:

show ipv6 route static

7. Show route information for a VLAN:

show ipv6 route vlan <1-4084>

Example

Switch:1#show ipv6 route count-summary

IPv6	6 Route Sum	mary						
	TOTAL	OSPF	RIP	BGP	STATIC	LOCAL		
	13	10	0	0	1		2	

Switch:1#show ipv6 route static

Static Route	e Inf	ormatio	on			
DEST-IP NEXT-HOP		IFIND FERENCI	•	BRT/'	TUN) ENABLE	STATUS
0:aa:1:0:0:0:0:0 NotReachable0:2910:0:0:0:0:0:10	64	0	(0)	enable	
0:aa:1:1:0:0:0:0 NotReachable0:2911:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:2:0:0:0:0 NotReachable0:2912:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:3:0:0:0:0 NotReachable0:2913:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:4:0:0:0:0 NotReachable0:2914:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:5:0:0:0:0 NotReachable0:2915:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:6:0:0:0:0 NotReachable0:2916:0:0:0:0:0:10	64	0	(0 5)	enable	
0:aa:1:7:0:0:0:0	64	0	(0)	enable	

NotReachable0:2917:0:0:0:0:0:0 5 8 out of 42 Total Num of Static Routes displayed. Global IPv6 Static Routes Admin Status: enable

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, BGP, 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.
{slot/port}	Identifies a single slot and port.
static	Shows static IPv6 routes.
vlan<1-4084>	Shows route entries for a specific VLAN ID.

IPv6 basic configuration using EDM

Configuring IPv6 forwarding

Enable the device to forward IPv6 traffic. By default, IPv6 forwarding is disabled, which means you can only use local IPv6 connections, and traffic does not traverse an IPv6 network.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click IPv6.
- 3. Click the **Globals** tab.
- 4. Select forwarding.
- 5. Click Apply.

Globals field descriptions

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

Name	Description
Forwarding	Configures whether this entity is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, this entity. Select forwarding to act as a router. Select notForwarding to not act as a router. The default is notForwarding.
	You must enable forwarding to terminate IPv6 packets on the CP, for example, 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.
IcmpRedirectMsg	Enables ICMP redirect messages. The default is disabled.
IcmpErrorInterval	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000. An entry of 0 seconds results in no sent ICMPv6 error messages
IcmpErrorQuota	Configures the number of ICMP error messages that can be sent during the ICMP error interval. A value of zero specifies not to send any. The default value is 50.
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.

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 *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500 and *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000,* NN46250-503.

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 following folders: **Configuration > IPv6**.
- 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.			
Vlanld	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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. The system has 1536 MAC addresses. The last four addresses are reserved.
	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.
ReversePathCheckEnable	Enables IPv6 reverse path checking.
	The default is disabled.
ReversePathCheckMode	Specifies the reverse path checking mode as one of the following:
	 exist-only— in this mode reverse path checking checks to determine whether the source address for the incoming packet exists in the routing table. If the source entry is found, the packet is forwarded; if not, the packet is discarded
	 strict—in this mode reverse path checking checks to determine whether the source address for the incoming packet exists in the routing table. If the source entry is found, reverse path checking

Name	Description
	checks further to determine whether the source interface address matches the packet incoming interface address. If they match, the packet is forwarded; if not, the packet is discarded.
ForwardingEnabled	Indicates whether IPv6 forwarding is enabled.
	The default is disabled.
RSMLTEnable	Shows whether RSMLT is enabled on the 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.

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. This procedure uses the main IPv6 navigation path where you can assign IPv6 addresses to both types of interfaces.

Procedure

- 1. In the navigation tree, 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 one of the following:
 - port
 - VLAN
- 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.			
	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.			
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.			

Creating IPv6 static routes

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

Before you begin

• Enable IPv6 forwarding.

Procedure

- 1. In the navigation tree, expand the following folders: Configuration > IPv6.
- 2. 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.
- 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.

Name	Description
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.

Viewing route information

View routes information to view the current configuration.

Procedure

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

3. Click 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. The range is 0 to 128.
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 <u>Globals field descriptions</u> on page 108.

Chapter 4: Neighbor discovery

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

Before you begin

Avaya includes the IPv6 neighbor discovery feature in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

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.

ND messages

The following table compares the ICMP message types.

Table 6: 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 switches can adjust their MAC tables.
Router solicitation message (optional)	Router solicitation message (required)	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

IPv4 function	IPv6 function	Description
		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 7: 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.
	🔁 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

State	Description
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.

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 VSP 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:

1. 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 ACLI

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[-slot/port][,...]} OF interface
vlan <1-4084>
```

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 prefixinterface 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.

Variable	Value
	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.
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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

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[-slot/port][,...]} or interface
vlan <1-4084>
```

- 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.

Variable	Value
	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>	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 0.
rtr-advert-max-interval <4–1800>	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface.
	The default is 600.
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.

Variable	Value
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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

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:

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

VSP-switch:1(config-if) # show ipv6 nd interface GigabitEthernet 5/1

====:			 =======	Por	====== t Ipv6	Nd		======			
IFID	BTR	RTR- ADV	 MIN- INT FLAG	LIFE- TIME CONF	MANAG	OTHER LIMIT	DAD-NS TIME			REACH- ABLE	RETRANS- MIT
			200 d) = Defa				1	1300(i)	64 (d)	0(d)	0(d)

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:

VSP-switch:1(config-if) #ipv6 interface reachable-time 60000

VSP-switch:1(config-if) #ipv6 nd reachable-time 0

VSP-switch:1(config-if)#show ipv6 nd interface GigabitEthernet 5/1

		 	 Port	====== t Ipv6	Nd	 			
IFID	BTR	 MAX- INT	 LIFETIME CONF FLAG	MANAG		 MTU TIME		REACH- ABLE	RETRANS- MIT
			0 Default,			1500(i)	64 (d)	0(s)	0 (d)

VSP-switch:1(config-if)#show ipv6 nd interface GigabitEthernet 1/1

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.

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} mac <0x00:0x00:0x00:0x00:0x00> [vlan <1-4084>]

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 10/11 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}	Identifies a single slot and port.
vlan <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the destination cache for all interfaces:

show ipv6 dcache

3. View the destination cache for a brouter port:

show ipv6 dcache gigabitethernet {slot/port}

4. View the destination cache for a management port:

show ipv6 dcache mgmtethernet {slot/port}

5. View the destination cache for a VLAN:

show ipv6 dcache vlan <1-4084>

6. Clear the destination cache:

```
clear ipv6 dcache [gigabitethernet <slot/port>][mgmtethernet <slot/
port>][vlan <1-4084>]
```

Example

Switch:1(config-if)#	#show ipv6 dcach	1e			
	IPv6 Destinat	ion Cache Info	rmation		
Destination Address	NEXT HOP	VID/BID/TID	IF_TYPE	IF_DATA	PMTU PMTU_

2:0:0:0:0:0:0:0:36 3:0:0:0:0:0:0:0:36 4:0:0:0:0:0:0:0:36	0:0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0:0	V-2 V-3	real		 1500	0
3:0:0:0:0:0:0:36	0:0:0:0:0:0:0:0:0			-	1500	\cap
		17_2				
4:0:0:0:0:0:0:36			real	-	1500	0
	0:0:0:0:0:0:0:0:0	V-4	real	-	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	6/7	real	-	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-2	real	-	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-3	real	-	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-4	real	-	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	T-25	real	-	1280	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-2	virtual	rsmlt	1500	0
ff02:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-3	virtual	vrId-1	1500	0
ff02:0:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-4	virtual	vrId-1	1500	0
ff02:0:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-4	virtual	vrId-10	1500	0
ff02:0:0:0:0:0:0:0:1	0:0:0:0:0:0:0:0:0	V-3	virtual	vrId-255	1500	0
ff02:0:0:0:0:0:0:12	2 0:0:0:0:0:0:0:0:0	V-3	virtual	vrId-1	1500	0
ff02:0:0:0:0:0:0:12	2 0:0:0:0:0:0:0:0:0	V-4	virtual	vrId-1	1500	0
ff02:0:0:0:0:0:0:12	2 0:0:0:0:0:0:0:0	V-4	virtual	vrId-10	1500	0
ff02:0:0:0:0:0:0:12	2 0:0:0:0:0:0:0:0	V-3	virtual	vrId-255	1500	0
ff02:0:0:0:0:0:0:0	5 0:0:0:0:0:0:0:0	V-2	real	_	1500	0
ff02:0:0:0:0:0:0:0:16	5 0:0:0:0:0:0:0:0	V-3	real	_	1500	0
ff02:0:0:0:0:0:0:0:16	5 0:0:0:0:0:0:0:0:0	V-4	real	_	1500	0
ff02:0:0:0:0:0:0:0		T-25	real	_	1280	0
ff02:0:0:0:0:0:0:0		V-2	virtual	rsmlt	1500	0
ff02:0:0:0:0:0:0:16		V-3	virtual	vrId-1	1500	0
ff02:0:0:0:0:0:0:0:16		V-4	virtual	vrId-1	1500	Õ
ff02:0:0:0:0:0:0:16		V-4	virtual	vrId-10	1500	0
ff02:0:0:0:0:0:0:0:16		V-3	virtual	vrId-255		0

Variable definitions

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

Variable	Value
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port}	Identifies a single slot and port.

Neighbor discovery configuration using EDM

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 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.

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.
VLanld	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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.

Name	Description
	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.
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.

Name	Description
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).

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.
- 3. 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.

Name	Description
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.
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.

Name	Description
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.

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 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, in the range of 0-65535.
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

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.

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 (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 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

This chapter provides concepts and procedures to complete IPv6 DHCP Relay configuration.

Before you begin

Avaya includes the IPv6 DHCP Relay feature in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

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 <u>Host autoconfiguration</u> on page 61.

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. This release 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 VSP 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 servers to which a relay can send a message from one client is 10.
- You can configure a maximum of 512 forwarding paths per system.

DHCP Relay configuration using ACLI

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.

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.

You can configure a maximum of 512 forwarding paths.

Procedure

1. Enter Global Configuration mode:

enable configure terminal

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[-slot/port][,...]}
```

OR

interface vlan <1-4084>

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.

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 5/1
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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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[-slot/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).
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[-slot/port][,...]} OF interface
vlan <1-4084>
```

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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

Viewing DHCP Relay information

View DHCP Relay information to display the current configuration for the forwarding path and the interface configuration.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View forwarding path information:

show ipv6 dhcp-relay fwd-path

3. View IPv6 DHCP Relay interface configuration:

```
show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[-slot/
port][,...]}|vlan <1-4084>}
```

😵 Note:

The **no ipv6 dhcp-relay** command disables DHCP on the interface but does not delete the entry.

Example

Switch:1(config-if)#show ipv6 dhcp-relay fwd-path

DHCPv6 Fwd-path					
INTERFACE			SERVER		ENABLE
1111:0:0:0:0:0:0:1111			1234:0:0:	0:0:0:1234	enable
Switch:1(config-if)#show ipv6 dhcp-relay interface gigabitEthernet 5/1					
Port Dhcpv6					
PORT NUM	IF INDEX	MAX HOP	DHCP-RELAY	REMOTE ID	
5/1	320	30	enable	disable	

Variable definitions

Use the information in the following table to help you use the **show ipv6 dhcp-relay** command.

Variable	Value
{slot/port[-slot/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).
vlan<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

DHCP Relay configuration 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.

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.

You can configure a maximum of 512 forwarding paths.

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.

Name	Description
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.

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 and the **DHCP Relay** tab for brouter ports.

Name	Description
lfIndex	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.
RemoteIdEnabled	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).

Name	Description
DhcpEnabled	Enables (true) or disables (false) DHCP Relay for an interface with an existing DHCP Relay configuration.
	This field appears only on the DHCP Relay tab for a brouter port if you modify an existing configuration. This field does not appear if you create a new DHCP Relay port configuration.

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.
- 3. 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** tab.

Name	Description
lfindex	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 tree, expand the following folders: **Configuration > Edit > Port**.
- 3. Click IPv6.
- 4. Click the DHCP tab.
- 5. Double-click a cell to change the value.
- 6. Click Apply.

Interface field descriptions

Use the data in the following table to use the **Interface** tab and the **DHCP Relay** tab for brouter ports.

Name	Description
lfindex	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 only on the DHCP Relay tab for a brouter port if you modify an existing configuration. This field does not appear if you create a new DHCP Relay port configuration.

Chapter 6: OSPFv3

This chapter provides concepts and procedures to complete IPv6 Open Shortest Path First (OSPF)v3 configuration.

Before you begin

Avaya includes the IPv6 OSPFv3 feature in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

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).

In IPv6 the term link replaces the IPv4 terms subnet and network. 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.

😵 Note:

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 pariticpate 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, <u>Configuring OSPF on a port or VLAN</u> on page 99.

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.

Deprecation of MOSPF for IPV6

OSPFv3 in RFC 5340 deprecates Multicast Extensions to OSPF (MOSPF) support, and its attendant protocol fields. The Infinity IPv6 OSPF module has some MOSPF code that is originally ported from phase2 IPv4 OSPF, but never used and fully implemented.

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.

OSPFv3 configuration using ACLI

Configuring OSPF globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Procedure

1. Enter Global Configuration mode:

enable

```
configure terminal
```

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 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.

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 inter- area 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.

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). 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.

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.
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 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[-slot/port][,...]} or interface
vlan <1-4084>
```

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.

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. A value of zero indicates the metric value depends on the speed of the interface, when the state of the interface is up.
	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 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.

Variable	Value
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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

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

3. View OSPF areas:

show ipv6 ospf area

4. View OSPF interface information

```
show ipv6 ospf interface [gigabitEthernet {slot/port}|vlan <1-4084>]
```

5. View OSPF interface timers:

show ipv6 ospf int-timers

6. View the link-state database (LSDB) table:

```
show ipv6 ospf lsdb [adv-rtr <A.B.C.D>] [area <A.B.C.D>] [interface
fastEthernet|gigabitEthernet {slot/port}|vlan <1-4084>] [lsa-type
<1-8>] [lsid <0-4294967295>] [scope <1-3>][detail]
```

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

Example

Switch:1#show ipv6 ospf

	OSPFv3 Global In	IOTMATION ====================================		
router-id admin-state version area-bdr-rtr-state as-bdr-rtr-state as-scope-lsa-count lsa-checksum originate-new-lsas rx-new-lsas ext-lsa-count Helper mode	: DI : 3 : FA : FA : 0 : 0 : 0 : 0 : 0 : 0			
Switch:1>show ipv6 ospf ar	ea			
	OSPF Are	a		
AREA_ID STUB_AREA	NSSA IMPORT_SUM			
0.0.0.0 false STUB_METRIC SPF_RUNS BDR_R	false true TR_CNT ASBDR_RTR_	always CNT LSA_CNT :	LSACK_SUM	
10 0 0	0	0	0	
Switch:1#show ipv6 ospf interface Total ospf areas: 1 Total ospf interfaces: 2				
	OSPF Inter	face		
IFINDX(VID/BRT) AREAID				IFTYPE
331 (5/12) 0.0.0.0	ena DOWN		0.0.0.0 0.0.0.0	PT-PT 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 Switch:1#show ipv6 ospf neighbor ______ _____ ____ OSPF Neighbor _____ IFINDX (VID/BRT) NBRROUTERID NBRIPADDR STATE TTL _____ (10/19) 97.146.128.0 fe80:0:0:2ef4:c5ff:fe92:8a00 Restart 120 331 1 out of 1 Total Num of Neighbor Entries displayed. _____ OSPF Virtual Neighbor NBRAREAID NBRROUTERID VIRTINTFID NBRIPV6ADDR STATE _____ 0 out of 0 Total Num of Virtual Neighbor Entries displayed. _____ OSPF NBMA Neighbor _____ INTERFACE NBRROUTERID NBRIPADDR STATE _____ 0 out of 0 Total Num of NBMA Neighbor Entries displayed. OSPF NBMA Neighbor _____ INTERFACE NBRROUTERID NBRIPADDR 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 lsdb 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.
{slot/port}	Identifies a single slot and port.
Isa-type <1-8>	Shows information for the specified LSA type.
lsid <0-4294967295>	Shows information for the specified link-state ID.

Variable	Value
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.
vlan <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 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[-slot/port][,...]} OF interface
vlan <1-4084>
```

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 fe80:0:0:0:8217:7dff: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.

Enabling OSPF route redistribution

Enable redistribution to announce routes, of a certain source protocol type, into the OSPFv3 domain.

You can redistribute directly connected routes, IPv6 static routes, and IPv6 BGP routes into OSPFv3.

About this task

If you do not configure a redistribution entry, the switch generates external LSAs for non-OSPF routes.

You can also redistribute directly connected routes, IPv6 static routes, and IPv6 OSPFv3 routes into BGP.

For information about BGP for IPv6, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Enable IPv6 router redistribution:

ipv6 redistribute <bgp|direct|static> enable

Example

Announce IPv6 static routes into the OSPFv3 domain:

Switch:1(config-ospf)#ipv6 redistribute static enable

Variable definitions

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

Variable	Value
<bgp direct static></bgp direct static>	Shows the source protocol from which to receive routes to insert into the OSPFv3 domain. The possibilities are direct routes, static routes, and BGP routes. By default, no routes are announced. Route redistribution is disabled.

Viewing the status of OSPFv3 redistribution

View the status of OSPFv3 route redistribution to verify the current configuration. You can redistribute directly connected routes, IPv6 static routes, and IPv6 BGP routes into OSPFv3.

For information about IPv6 route redistribution into BGP, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the current configuration:

show ipv6 ospf redistribute

Example

```
Switch:1#show ipv6 ospf redistribute
```

OSPF Redistribute List		
direct static bgp	: disabled : disabled : enabled	

OSPFv3 configuration using EDM

Configuring OSPF globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 3. Click the Globals tab.
- 4. Type the router ID, in the format of an IPv4 address.
- 5. Select enabled.
- 6. 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.

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.

Name	Description
ASBdrRtrStatus	Configures the router as an autonomous system boundary router. The default is disabled (clear).
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.

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 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.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 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
ld	Specifies a 32–bit integer to uniquely identify an area. Use 0.0.00 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. The possible values are always or candidate. The default is candidate.

Name	Description
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 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.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 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.

nssaExternalLink 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
ArealD	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. nssaExternalLink 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 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.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 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.

Name	Description
	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.

Creating an OSPF interface

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 124. To configure OSPFv3 on an IPv6 VLAN, see <u>Creating an OSPF VLAN interface</u> on page 116.

Before you begin

• The IPv6 interface must exist.

Procedure

Configure OSPF on an IPv6 port:

- a. In the Device Physical view, select a port.
- b. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
- c. Click IPv6.
- d. Click the IPv6 OSPF Interface tab.
- e. Click Insert.
- f. Select the area ID.
- g. Select enabled.
- h. Click Insert.

IPv6 OSPF Interface field descriptions

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
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.

Use the data in the following table to use the IPv6 OSPF Interface tab.

Name	Description
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. A value of zero indicates the metric value depends on the speed of the interface, when the state of the interface is up. The default value for a brouter port or VLAN is 1.

Creating 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 124.

Before you begin

• The IPv6 interface must exist.

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 IPv6.

- 6. Click the IPv6 OSPF Interface tab.
- 7. Click Insert.
- 8. Select the area ID.
- 9. Select enabled.
- 10. 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.
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
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

Name	Description
	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. A value of zero indicates the metric value depends on the speed of the interface, when the state of the interface is up. The default value for a brouter port or VLAN is 1.

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 ASexternal-LSAs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.

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 OSPF area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IPv6**.
- 2. Click OSPF.
- 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.
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.
Routerld	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.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 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.
Routerld	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

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 **OSPF**.
- 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
lfIndex	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

Enabling OSPF route redistribution

Enable redistribution to announce routes of a certain source protocol type into the OSPFv3 domain.

You can redistribute directly connected routes, IPv6 static routes, and IPv6 BGP routes into OSPFv3.

About this task

If you do not configure a redistribution entry, the switch generates external LSAs for non-OSPF routes.

You can also redistribute directly connected routes, IPv6 static routes, and IPv6 OSPFv3 routes into BGP.

For information about BGP for IPv6, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
- 2. Click **OSPF**.
- 3. Click the **Redistribute** tab.
- 4. For the type of route source, double-click the cell in the **Enable** column to change the value.
- 5. Select enable.
- 6. Click Apply.

Redistribute field descriptions

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

Name	Description
DstVrfld	Shows the ID of the destination virtual router and forwarder (VRF). Because IPv6 is not virtualized, the value is 0 for the Global Router.
Protocol	Shows the routing protocol that receives the external routing information. In this case, the routing protocol is OSPFv3.
SrcVrfld	Shows the ID of the source VRF. Because IPv6 is not virtualized, the value is 0 for the Global Router.
RouteSource	Shows the source protocol from which to receive routes to insert into the OSPFv3 domain. The possibilities are direct routes, static routes, and BGP routes.
Enable	Configures the status of route redistribution. The default is disable.

Modifying an OSPF interface

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface. An IPv6 interface can be a port, or VLAN.

Before you begin

• The OSPF interface must exist.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IPv6**.
- 2. Click OSPF.
- 3. Click the Interfaces tab.
- 4. Double-click a cell to edit the value.
- 5. Click Apply.

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
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

Name	Description
	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.
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. A value of zero indicates the metric value depends on the speed of the interface, when the state of the interface is up. The default value for a brouter port or VLAN is 1.

Viewing OSPF neighbors

View OSPF neighbors to see routers with interfaces to a common network.

The OSPF 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 OSPF.
- 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

Name	Description
	• full
Nbrifid	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 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 OSPF backbone.

Procedure

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

Virtual Neighbors field descriptions

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

Name	Description
Area	Shows the ID for the transit area.
Rtrld	Shows the ID for the neighboring router in the autonomous system.
Locallfindex	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

Name	Description
	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.
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 7: VRRP

This chapter provides concepts and procedures to complete IPv6 Virtual Router Redundancy Protocol (VRRP) configuration.

Before you begin

Avaya includes the IPv6 VRRP feature in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

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 Avaya 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.

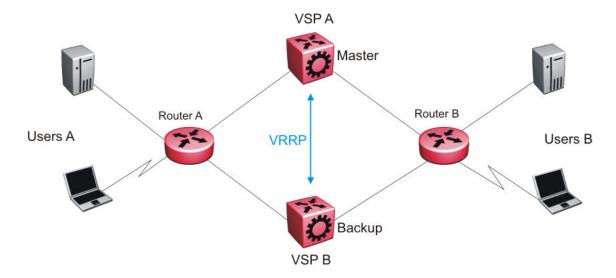
VRRP

With the current implementation of VRRP, 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.

😵 Note:

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 Avaya 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.

Avaya recommends that you 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 the interswitch trunk (IST) link toward the master VRRP router. In this case, the IST link 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 Layer 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, Avaya provides the fast advertisement interval.

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 Avaya products with the same configuration.

IPv6 VRRP and ICMP redirects

In IPv6 networks, do not enable ICMP redirects on VRRP VLANs. If you enable this option (using the ipv6 icmp redirect-msg command), VRRP cannot function. The option is disabled by default.

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 interswitch trunk (IST) 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.

VRRP configuration using ACLI

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 back up.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]} OF interface
vlan <1-4084>
```

- 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 <1-255> 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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[-slot/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).

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]
```

3. View the configuration information for one or more ports:

```
show ipv6 vrrp interface gigabitethernet [{slot/port[-slot/port]
[,...]}] [verbose]
```

4. View the configuration information for one or more VLANs:

show ipv6 vrrp interface vlan [<1-4084>] [verbose]

- 5. View the configuration information for one or more virtual router IDs: show ipv6 vrrp interface vrid <1-255> [verbose]
- 6. View VRRP backup address information:

show ipv6 vrrp address

7. View VRRP backup address information for a link-local address:

show ipv6 vrrp address link-local WORD<0-127> [verbose]

8. View VRRP backup address information for a virtual router ID:

show ipv6 vrrp address vrid <1-255>

Example

Switch:1>show ipv6 vrrp address

			VRRP Ir		GlobalRoute				
VRID	P/V	IP			MAC		STA	TE C	ONTROL
12	5/1	fe80:0:0	:0:0:0:0:1234		00:00:5e	:00:02:0c	Ini	t D	isabled
VRID	P/V	MASTER			PRIO AD	V UP TIM	E		
12	5/1	0:0:0:0:	0:0:0:0		100 1	0 day(s), 0	0:00:00	
VRID	P/V	CRITICAL	IP			CRITICAL ENABLED		ACCEPT MODE	I
12	5/1	0:0:0:0:	0:0:0:0			No		disabl	e
VRID	P/V		BACKUP-MASTER STATE		(ENABLED)	ACTION	HLD DWN	REM	[
12	5/1	disable	down	400	(YES)	none	30	0	
VRID	P/V	GLOBAL A	DDRESS						
 12	5/1	1111::22	22/64						

Flags Legend:

HLD DWN: Configured hold-down timer value, REM: REMaining hold-down timer value

--More-- (q = quit)

Variable definitions

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

Variable	Value
{slot/port[-slot/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).
link-local WORD<0-127>	Displays information by link-local IPv6 address.
verbose	Displays extended information.
vlan [<1-4084>]	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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.

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[-slot/port][,...]} or interface
vlan <1-4084>
```

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 <1-255> fast-adv enable [fast-adv-int <200-1000>]
```

8. Configure the holddown timer:

ipv6 vrrp <1-255> holddown-timer <0-21600>

9. Configure the priority for the VRRP router:

ipv6 vrrp <1-255> priority <1-255>

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.
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 IST link. 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.

Variable	Value
critical-ipv6-addr <i>WORD<0-46></i>	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.

VRRP configuration using EDM

Configuring 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.

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 VSP does respond to SNMP get requests to the physical IP address.

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 pane, expand the following folders: **Configuration > IPv6**.
- 2. Click VRRP.
- 3. Click the **Interface** tab.
- 4. Click Insert.
- 5. Beside the **IfIndex** field, click **Port** or **VLAN**.
- 6. Select a port or VLAN.
- 7. Click OK.
- 8. Type the virtual router ID.
- 9. Type the primary IP address.
- 10. Click Insert.

Interface field descriptions

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

Name	Description
lfindex	Shows the index value that uniquely identifies the interface to which this entry applies.
InetAddrType	Specifies the address type for the VRRP interface. In this case, IPv6.
Vrld	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.
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

Name	Description
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. Only the master router sends advertisements.
MasterlpAddr	Specifies the IP address of the physical interface of the master virtual router that forwards packets sent to the virtual IP addresses associated with the virtual router.
UpTime	Shows the time interval, in hundredths of a second, since the virtual router was initialized.
CriticallpAddr	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.
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 IST link. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP switch traffic forwarding is enabled or disabled.
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. You must configure the same value on all participating routers. 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.

Name	Description	
Action	Lists options to override the holddown 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.	

Configuring VRRP for a VLAN

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.

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 does, however, respond to SNMP get requests to the physical IP address.

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 IPv6.
- 6. Click the VRRP tab.
- 7. Click Insert.
- 8. Type the Vrld.
- 9. Type the primary IP address.

10. Click Insert.

VRRP field descriptions

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

Name	Description		
lfindex	Shows the index value that uniquely identifies the interface to which this entry applies.		
InetAddrType	Specifies the address type for the VRRP interface. In this case, IPv6.		
Vrld	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.		
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.		
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. You must configure the same value on all participating routers. The default is 200.		
	Enter the values in multiples of 200 milliseconds.		
Advinterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second. Only the master router sends advertisements.		

Name	Description
Action	Lists options to override the holddown 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.
MasterlpAddr	Specifies the IP address of the physical interface of the master virtual router that forwards packets sent to the virtual IP addresses associated with the virtual router.
UpTime	Shows the time interval, in hundredths of a second, since the virtual router was initialized.
CriticallpAddr	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.
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 IST link. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP switch traffic forwarding is enabled or disabled.
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.
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.

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.

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 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.

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 AssociatedIPAddr.
- 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
lpAddr	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 8: RSMLT

This chapter provides concepts and procedures to complete IPv6 Routed Split MultiLink Trunking (RSMLT) configuration.

Before you begin

Avaya includes the IPv6 RSMLT feature in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

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 IST 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 IST 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 IST.

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:

You cannot clear the static IPv6 routes during the RSMLT holddown timer. The RSMLT holddown timer defines how long the recovering or restarting system remains in non-Layer 3 forwarding mode for the peer route MAC address. If you try to clear the static IPv6 routes during the holddown timer, the system displays the following output: Static routes cannot be cleared until the RSMLT holddown period is done (in 39 seconds). Try again later.

For more information about the IPv4 RSMLT model, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505. This section focuses on the differences between the IPv4 and IPv6 models.

For an example IPv6 RSMLT network topology, see *Network Design Reference for Avaya Virtual Services Platform 9000,* NN46250-200.

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 can not configure an IST 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 ACLI

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 <1-4084>
```

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 ip rsmlt 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.

Table continues...

Variable	Value
	Note:
	You cannot clear the static IPv6 routes during the RSMLT holddown timer, which defines how long the RSMLT device does not particpate in Layer 3 forwarding. If you try to clear the static IPv6 routes during the holddown timer, the system displays the following output: Static routes cannot be cleared until the RSMLT holddown period is done (in 39 seconds). Try again later.
holdup-timer <0-9999>	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 180. 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
<1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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	fo - Glo	balRou	ter	
			=======	======		
VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
101		00:24:7f:9e:da:01		-		
102	102.1.1.32	00:24:7f:9e:da:02	Enable	Up	60	180
VID	SMLT ID					
	101					
102	102					
VID	IPv6	MAC	ADMIN	OPER	HDTMR	HUTMR
101		00:24:7f:9e:da:01	Enable	Up	100	200
	1010:0:0:0:0:0:0: 1010:0:0:0:0:0:0:					
102	fe80:0:0:0:224:	:7fff:fe9e:da01/128 00:24:7f:9e:da:02	Frahlo	Un	60	1 0 0
102	1020:0:0:0:0:0:0:		Ellabre	op	00	100
	1020:0:0:0:0:0:0:	:0:32/64				
	fe80:0:0:0:224:	7fff:fe9e:da02/128				
VID	SMLT ID					
	101					
102	102					

			Ip Rsm =======	lt Peer Inf =======	o - Glob	alRout ======	er ======	
VID	IP		MAC		ADMIN	OPER	HDTMR	HUTMR
	101.1.1.33 102.1.1.33							
VID	HDT REMAIN	HUT	REMAIN	SMLT ID				
	60 60							
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 :33/64	f:9e:ea:01	Enable	Up	100	200
102	1020:0:0:0: 1020:0:0:0: fe80:0:0:0:	0:0:0	00:24:7 :0/64 :33/64	f:9e:ea:00	Enable	Up	60	180
VID	HDT REMAIN	HUT	REMAIN	SMLT ID				
101 101	60	180						
101 102 102	60	180						

RSMLT Peer Info: rsmlt-peer-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

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.
- 6. 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.

Table continues...

Name	Description		
	Note:		
	You cannot clear the static IPv6 routes during the RSMLT holddown timer period, which defines how long the RSMLT device does not participate in Layer 3 forwarding. If you try to clear the static IPv6 routes during the holddown timer period, the system displays the following output: Static routes cannot be cleared until the RSMLT holddown period is done (in 39 seconds). Try again later.		
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 180.		
	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.
- 2. 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

- 1. 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.
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 0.
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.
OperStatus	Displays the RSMLT operating status as either up or down.
Smitid	Specifies the ID range for the SMLT.
Vlanld	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
lfIndex	Shows the route SMLT operation index.
lpv6Addr	Configures the IPv6 address of the RSMLT interface.
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.
Vlanld	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. 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 9: 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

Before you begin

Avaya includes IPv6 support in the Base License.

For more information about feature licensing, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Viewing IPv6 connections using ACLI

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

3. Display IPv6 TCP listener information for the specified IPv6 address:

show ipv6 tcp listener

4. Display IPv6 TCP properties

show ipv6 tcp properties

5. Display IPv6 TCP statistics

show ipv6 tcp statistics

6. Display IPv6 UDP information:

show ipv6 udp endpoints

OR

show ipv6 udp local addr WORD<0-128> [slot/port]

OR

show ipv6 udp remote addr WORD<0-128> [slot/port]

Example

	now ipv6 tcp connectior			
TCP connect	ion table info			
LOCALPORT 23 33471	LOCALADDR 3000:0:0:0:0:0:0:0:1 3000:0:0:0:0:0:0:0:2 4000:0:0:0:0:0:0:0:1 4000:0:0:0:0:0:0:0:2	REMOTEPORT	REMOTEADDR	STATE established established
	now ipv6 tcp listener			
	er table info			
LOCALPORT 21 22 23 80 443 513	LOCALADDR 0:0:0:0:0:0:0 0:0:0:0:0:0:0 0:0:0:0:0:	D:0:0: D:0:0: D:0:0: D:0:0: D:0:0:		
	now ipv6 tcp properties o global prioerties cor			
RtoAlgorigh RtoMin RtoMax MaxConn	um constant 5002 millisecond 60128 millisecon 127			
	now ipv6 tcp statistics	5		
ActiveOpens	s: 0			

PassiveOpens:	6
	č
AttemptFails:	0
EstabResets:	0
CurrEstab:	2
InSegs:	50
OutSegs:	38
RetransSegs:	0
InErrs:	0
OutRsts:	2
HCInSegs:	38
HCOutSegs:	38

```
Switch:1>show ipv6 udp endpoints
```

	UDP endpoint table info		
LOCALPORT REMOTEPORT	LOCALADDR REMOTEADDR	INSTANCE	PROCESS
69 0	0:0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0	1219584048	0
161 0	0:0:0:0:0:0:0:0 0:0:0:0:0:0:0:0	1219585596	0

Viewing IPv6 connections using EDM

Viewing TCP global information

View TCP and UDP information to view the current configuration.

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

- In the navigation tree, expand the following folders: Configuration > IP or Configuration > IPv6.
- 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.

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 tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
- 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.

Table continues...

Name	Description
RemAddressType	Displays the type (IPv6 or IPv4) for the remote address of the TCP connection.
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.

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

- In the navigation tree, expand the following folders: Configuration > IP or Configuration > IPv6.
- 2. Double-click **TCP/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.

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 tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
- 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 10: IPv6 configuration examples

The following sections show configuration examples for IPv6 deployment options.

For more configuration examples that include configuration of additional Avaya products, see *IPv6* for VSP 9000 Technical Configuration Guide, NN48500–634.

OSPFv3

This section shows an example of OSPFv3 configuration. The following figure shows the network.

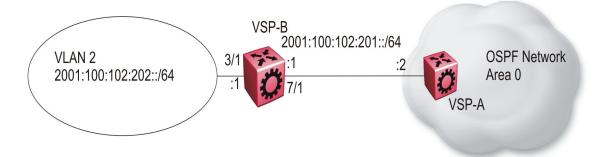


Figure 12: OSPFv3 configuration

To complete the configuration, you must perform the following actions:

- Configure an IPv6 VLAN (VLAN 2) with port member 3/1.
- Configure an IPv6 brouter port (7/1).
- Use IPv6 address 2001:100:102::/64.

Configure VLAN 2 and add port members.

```
vlan create 2 type port-mstprstp 1
vlan mlt 2 4
vlan members 2 3/1 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 7/1 with IPv6 and OSPFv3.

```
interface gigabitethernet 7/1
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:

VSP-switch:1#show ipv6 ospf area

			OSP	F Are	ea					
AREA_ID	======================================	NSSA	IMPOR	===== T_SUN	====== M TRANS	ROLE	 IRANS_	STATE		
0.0.0.0 STUB_METRIC STUB	false _METRIC_TYP								NT LSA	.ck_su
10 ospf ^v	V3Metric	0		0		0		0	0	
VSP-switch:1#show	w ipv6 inte	rface	vlan 3	2						
	Vlan	Ipv6]	Interf	===== ace =====						
IF VLAN PHYSICA INDX INDX ADDRES:					ABLE	SMIT		FORWARI ING JS ADMIN STATUS		RPC MODE
2070 2 00:24:7f: al:7a:06	enable up	ETHEF	R 1500	64	30000	1000	disab	ole en- able		
		 V]	lan Ip	==== v6 Ac	ddress					
IPV6 ADDRESS				==== VL2	====== AN-ID	TY	====== PE	ORIGIN	STAT	===== 'US
2001:100:102:202 fe80:0:0:0:224:7		.06			2 2 2			IANUAL JINKLAYER		

1 out of 2 Total Num of Interface Entries displayed. 2 out of 5 Total Num of Address Entries displayed.

Dual stack core

This section shows an IPv6 deployment that includes a dual stack core. The following figure shows the network.

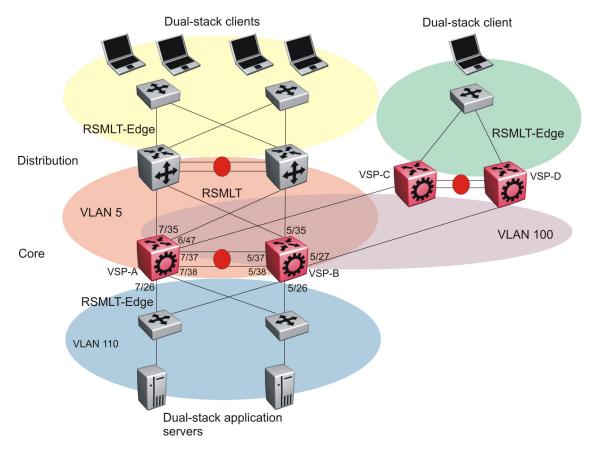


Figure 13: IPv6 deployment

The following example provisions the SMLT cluster VSP-A and VSP-B. The network core uses the following configuration:

- Two VSP 9000 core switches operate in an SMLT cluster.
- All edge switches are Layer 2 switches with an MLT link.
- The VSP 9000 core devices use RSMLT on the Layer 3 routing instances with the distribution devices.
- On VSP-A, the IST VLAN uses 10.1.2.1/30, and ports 7/37 and 7/38.
- On VSP-B, the IST VLAN uses 10.1.2.2/30, and ports 5/37 and 5/38.
- OSPFv2 and OSPFv3 run between the core and the distribution layer.
- The MLT ID between the core and distribution devices is 7.

- VSP-A uses the following addresses for VLAN 5 on port 7/35:
 - 10.1.5.1/24
 - fd12:0:0:1205::1/64
- VSP-B uses the following addresses for VLAN 5 on port 5/35:
 - 10.1.5.2/24
 - fd12:0:0:1205::2/64
- The MLT ID between the core and edge devices on VLAN 110 is 6. VLAN 110 is provisioned as RSMLT-Edge.

VSP-A configuration

```
# LACP CONFIGURATION
vlacp enable
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 7/37
loop-detect action mac-discard
exit
interface GigabitEthernet 7/38
loop-detect action mac-discard
# MLT CONFIGURATION
#
mlt 1 enable name "IST"
mlt 1 member 7/37-7/38
mlt 1 encapsulation dot1q
mlt 6 enable name "smlt6"
mlt 6 member 7/26
mlt 6 encapsulation dot1q
mlt 7 enable name "SMLT_to_Distribution"
mlt 7 member 7/35
mlt 7 encapsulation dot1q
mlt 8 enable
mlt 8 member 6/47
mlt 8 encapsulation dot1q
# VLAN CONFIGURATION
vlan create 2 name "IST" type port-mstprstp 1
vlan mlt 2 1
vlan members 2 7/37-7/38 portmember
interface Vlan 2
ip address 10.1.2.1 255.255.255.252 1
exit
vlan create 5 type port-mstprstp 1
vlan mlt 5 1
vlan mlt 5 7
vlan members 5 7/35,7/37-7/38 portmember
interface Vlan 5
```

```
ip address 10.1.5.1 255.255.255.0
ip ospf enable
ipv6 interface enable
ipv6 interface address fd12:0:0:1205:0:0:1/64
ip rsmlt
exit
vlan create 100 type port-mstprstp 1
vlan mlt 100 1
vlan mlt 100 8
vlan members 100 6/47 portmember
interface Vlan 100
ip address 10.1.5.3 255.255.255.0
ip ospf enable
ipv6 interface mac-offset 6
ipv6 interface enable
ipv6 interface address fd12:0:0:2015:0:0:0:1/64
ip rsmlt
exit
vlan create 110 type port-mstprstp 1
vlan mlt 110 1
vlan mlt 110 6
vlan members 110 7/26,7/37-7/38 portmember
interface Vlan 110
ip address 10.1.101.1 255.255.255.0 5
ip ospf enable
ipv6 interface enable
ipv6 interface address fd12:0:0:1201:0:0:0:1/64
ip rsmlt
ip rsmlt holdup-timer 9999
exit
# MLT INTERFACE CONFIGURATION
#
interface mlt 1
ist peer-ip 10.1.2.2 vlan 2
ist enable
exit
mlt 1 enable name "IST"
mlt 1 member 7/37-7/38
mlt 1 encapsulation dot1q
interface mlt 6
smlt
exit
interface mlt 7
smlt
exit
interface mlt 8
smlt
exit
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 6/47
default-vlan-id 100
no shutdown
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 7/26
default-vlan-id 110
no shutdown
```

```
vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 7/35
default-vlan-id 5
no shutdown
vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit.
interface GigabitEthernet 7/37
default-vlan-id 2
no shutdown
vlacp slow-periodic-time 10000 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 7/38
default-vlan-id 2
no shutdown
vlacp slow-periodic-time 10000 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ip address 1 10.1.1.1/255.255.255.255
ip ospf 1
  OSPF CONFIGURATION - GlobalRouter
#
router ospf enable
router ospf
router-id 10.1.1.1
# RSMLT CONFIGURATION
ip rsmlt edge-support
# IPV6 CONFIGURATION
ipv6 forwarding
# IPV6 OSPFV3 CONFIGURATION
router ospf ipv6-enable
router ospf
ipv6 router-id 10.1.1.1
```

```
exit
#
# IPV6 OSPF VLAN CONFIGURATION
#
interface vlan 5
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
exit
interface vlan 100
ipv6 ospf area 0.0.0.0
ipv6 ospf area 0.0.0.0
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
exit
```

VSP-B configuration

```
# LACP CONFIGURATION
#
vlacp enable
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 5/37
loop-detect action mac-discard
exit.
interface GigabitEthernet 5/38
loop-detect action mac-discard
# MLT CONFIGURATION
#
mlt 1 enable name "IST"
mlt 1 member 5/37-5/38
mlt 1 encapsulation dot1q
mlt 6 enable name "smlt6"
mlt 6 member 5/26
mlt 6 encapsulation dot1q
mlt 7 enable name "SMLT_to_Distribution"
mlt 7 member 5/35
mlt 7 encapsulation dot1q
mlt 8 enable
mlt 8 member 5/27
mlt 8 encapsulation dot1q
# VLAN CONFIGURATION
vlan create 2 name "IST" type port-mstprstp 1
vlan mlt 2 1
vlan members 2 5/37-5/38 portmember
interface Vlan 2
ip address 10.1.2.2 255.255.255.252 0
exit
vlan create 5 type port-mstprstp 1
vlan mlt 5 1
vlan mlt 5 7
vlan members 5 5/35,5/37-5/38 portmember
```

```
interface Vlan 5
ip address 10.1.5.2 255.255.255.0
ip ospf enable
ipv6 interface enable
ipv6 interface address fd12:0:0:1205:0:0:0:2/64
ip rsmlt
exit
vlan create 100 type port-mstprstp 1
vlan mlt 100 1
vlan mlt 100 8
vlan members 100 5/27,5/37-5/38 portmember
interface vlan 100
ip address 10.1.5.4 255.255.255.0
ip ospf enable
ipv6 interface mac-offset 5
ipv6 interface enable
ipv6 interface address fd12:0:0:2015:0:0:2/64
ip rsmlt
exit
vlan create 110 type port-mstprstp 1
vlan mlt 110 1
vlan mlt 110 6
vlan members 110 5/26,5/37-5/38 portmember
interface Vlan 110
ip address 10.1.101.2 255.255.255.0 1
ip ospf enable
ipv6 interface enable
ipv6 interface address fd12:0:0:1201:0:0:2/64
ip rsmlt
ip rsmlt holdup-timer 9999
exit
# MLT INTERFACE CONFIGURATION
#
interface mlt 1
ist peer-ip 10.1.2.1 vlan 2
ist enable
exit
interface mlt 6
smlt
exit
interface mlt 7
smlt
exit.
interface mlt 8
smlt
exit
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 5/26
default-vlan-id 110
no shutdown
vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 5/27
default-vlan-id 100
no shutdown
```

```
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 5/35
default-vlan-id 5
no shutdown
vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 5/37
default-vlan-id 2
no shutdown
vlacp slow-periodic-time 10000 timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
interface GigabitEthernet 5/38
default-vlan-id 2
no shutdown
vlacp slow-periodic-time 10000 timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
vlacp enable
no spanning-tree mstp force-port-state enable
no spanning-tree mstp msti 1 force-port-state enable
exit
# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
interface loopback 1
ip address 1 10.1.1.2/255.255.255.255
ip ospf 1
#
  OSPF CONFIGURATION - GlobalRouter
router ospf enable
router ospf
router-id 10.1.1.2
# RSMLT CONFIGURATION
ip rsmlt edge-support
# IPV6 CONFIGURATION
ipv6 forwarding
# IPV6 OSPFV3 CONFIGURATION
router ospf ipv6-enable
router ospf
ipv6 router-id 10.1.1.2
exit
```

#

```
# IPV6 OSPF VLAN CONFIGURATION
#
interface vlan 5
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
exit
interface vlan 100
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
interface vlan 110
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
exit
```

Verification

VSP-A:1#show ip interface

IP Interface - GlobalRouter							
INTERFACE	IP ADDRESS	NET MASK	BCASTADDR FORMAT	REASM MAXSIZE	VLAN ID	BROUTER PORT	
Clip1 Vlan2 Vlan5 Vlan100 Vlan110	10.1.1.1 10.1.2.1 10.1.5.1 10.1.5.3 10.1.101.1	255.255.255.255 255.255.255.255 255.255.		1500 1500 1500 1500 1500 1500	 2 5 100 110	false false false false false false	

VSP-B:1#show ip interface

IP Interface - GlobalRouter

INTERFACE	IP ADDRESS	======================================	BCASTADDR FORMAT	========= REASM MAXSIZE	====== VLAN ID	======== BROUTER PORT
Clip1 Vlan2 Vlan5 Vlan100 Vlan110	10.1.1.2 10.1.2.2 10.1.5.2 10.1.5.4 10.1.101.2	255.255.255.255 255.255.255.255 255.255.		1500 1500 1500 1500 1500	 2 5 100 110	false false false false false

VSP-A:1#show ipv6 forwarding Global forwarding

VSP-B:1#show ipv6 forwarding Global forwarding

: enable

: enable

VSP-A:1#show ipv6 address interface

	Address Information			
IPV6 ADDRESS	VID/BID/TID	======== TYPE	ORIGIN	STATUS
fd12:0:0:1205:0:0:0:1 fe80:0:0:0:21b:4fff:fe60:fa04 fd12:0:0:2015:0:0:0:1 fe80:0:0:0:21b:4fff:fe60:fa06 fd12:0:0:1201:0:0:0:1 fe80:0:0:0:21b:4fff:fe60:fa05	V-5 V-5 V-100 V-100 V-110 V-110 V-110	UNICAST UNICAST UNICAST	LINKLAYER MANUAL LINKLAYER	PREFERRED PREFERRED PREFERRED

VSP-B:1#show ipv6 address interface

Address Information

IPV6 ADDRESS	VID/BID/TID	======= TYPE	ORIGIN	STATUS
fd12:0:0:1205:0:0:0:2 fe80:0:0:0:21b:4fff:fe5f:fa03 fd12:0:0:2015:0:0:0:2 fe80:0:0:0:21b:4fff:fe5f:fa05 fd12:0:0:1201:0:0:0:2 fe80:0:0:0:21b:4fff:fe5f:fa01	V-5 V-5 V-100 V-100 V-110 V-110 V-110	UNICAST UNICAST UNICAST	LINKLAYER MANUAL LINKLAYER	PREFERRED PREFERRED PREFERRED

VSP-A:1#show ipv6 ospf interface

OSPE Interface

				0.51	f inceria	ace			
IFINDX	(VID/BRI	:)	AREAID	ADM	IFSTATE	METRIC	PRI	DR/BDR	IFTYPE
2053	(5)	0.0.0.0	ena	DR_OTHER	1	_	10.1.1.2 10.1.1.4	BROADCAST
2098	(100)	0.0.0.0	ena	DR	1		10.1.1.1 10.1.1.2	BROADCAST
2158	(110)	0.0.0.0	ena	DR	1	1	10.1.1.1 10.1.1.2	BROADCAST

VSP-B:1#show ipv6 ospf interface

OSPF Interface

	IFINDX	(VID/BRI	С)	AREAID	ADM	IFSTATE	METRIC	PRI	DR/BDR	IFTYPE
,	2053	(5)	0.0.0.0	ena	DR	1	1	10.1.1.2 10.1.1.4	BROADCAST
	2098	(100)	0.0.0.0	ena	BDR	1	1	10.1.1.1 10.1.1.2	BROADCAST
	2158	(110)	0.0.0.0	ena	BDR	1	1	10.1.1.1 10.1.1.2	BROADCAST

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 8: 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

Table continues...

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	41 Inverse neighbor discovery Solicitation Message		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.
Extended Unique Identifier (EUI)	A 64-bit format used in assigning addresses automatically to IPv6 interfaces.
IPv6 address owner	The Virtual Router Redundancy Protocol (VRRP) router that uses the link- local IPv6 address of the virtual router as the real interface link-local address. This router responds to packets addressed to the IPv6 address.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.
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.