



Configuration — OSPF and RIP

Avaya Virtual Services Platform 9000

3.0
NN46250-506, 01.02
January 2011

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 generally makes available to users of its products. 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 Web sites 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 its Hardware and Software ("Product(s)"). 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 Web site: <http://support.avaya.com>. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](http://support.avaya.com/licenseinfo/) ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. 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 AUTHORIZED RESELLER; 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").

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or Hardware provided by Avaya. All content on this site, the documentation 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

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://support.avaya.com/Copyright>.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation 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 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, and "Linux" is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: <http://support.avaya.com>.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://support.avaya.com>.

Contents

Chapter 1: New in this release.....	7
Features.....	7
Other changes.....	7
Chapter 2: Introduction.....	9
Chapter 3: Routing fundamentals.....	11
Routing protocols.....	11
IP addresses.....	12
VLANs and routing.....	16
Static routes.....	17
Black hole static routes.....	17
Circuitless IP.....	18
Route policies.....	19
IP routing features and considerations.....	23
Chapter 4: OSPF fundamentals.....	27
OSPF overview.....	27
Dijkstras algorithm.....	28
Autonomous system and areas.....	29
OSPF neighbors.....	32
Router types.....	33
OSPF interfaces.....	34
OSPF and IP.....	39
OSPF packets.....	39
Intra-area link-state advertisements.....	40
ASE routes.....	41
OSPF virtual links.....	41
OSPF ASBRs.....	42
OSPF metrics.....	44
OSPF security mechanisms.....	44
OSPF and route redistribution.....	45
OSPF configuration considerations.....	46
OSPF host route advertisements and nonbackbone areas.....	46
OSPF with switch clustering.....	46
Chapter 5: RIP fundamentals.....	49
Chapter 6: OSPF configuration using EDM.....	53
Configuring OSPF globally.....	54
Enabling OSPF globally.....	56
Configuring global default metrics.....	56
Configuring an OSPF interface.....	57
Changing an OSPF interface type.....	60
Configuring NBMA interface neighbors.....	61
Configuring OSPF interface metrics.....	62
Viewing all OSPF-enabled interfaces.....	63
Configuring OSPF on a port.....	64
Configuring OSPF on a VLAN.....	66

Creating stubby or not-so-stubby OSPF areas.....	69
Configuring stub area metrics advertised by an ABR.....	70
Inserting OSPF area aggregate ranges.....	71
Enabling automatic virtual links.....	73
Configuring a manual virtual interface.....	73
Viewing virtual neighbors.....	75
Configuring host routes.....	76
Enabling ASBR status.....	77
Managing OSPF neighbors.....	78
Viewing the link-state database.....	78
Viewing the external link-state database.....	79
Configuring route redistribution to OSPF.....	80
Forcing shortest-path calculation updates.....	82
Chapter 7: OSPF configuration using ACLI.....	83
Configuring OSPF globally.....	83
Configuring OSPF for a port or VLAN.....	86
Viewing OSPF errors on a port.....	88
Configuring OSPF areas on the router.....	89
Configuring OSPF aggregate area ranges on the router.....	90
Enabling automatic virtual links.....	92
Configuring an OSPF area virtual interface.....	92
Configuring an OSPF area on a VLAN or port.....	94
Configuring an OSPF host route.....	96
Configuring OSPF NBMA neighbors.....	97
Applying OSPF route acceptance policies.....	98
Viewing the OSPF link-state database.....	100
Viewing the OSPF external link-state database.....	100
Configuring route redistribution to OSPF.....	101
Configuring interVRF route redistribution for OSPF.....	103
Forcing shortest-path calculation updates.....	104
Chapter 8: RIP configuration using EDM.....	107
Configuring RIP globally.....	107
Configuring RIP interface compatibility.....	108
Configuring RIP on an interface.....	111
Configuring RIP on a port.....	113
Configuring RIP on a VLAN.....	115
Configuring route redistribution to RIP.....	117
Chapter 9: RIP configuration using ACLI.....	119
Configuring RIP globally.....	119
Configuring RIP on an interface.....	121
Example of configuring RIP on an interface.....	123
Configuring route redistribution to RIP.....	124
Configuring interVRF route redistribution for RIP.....	125
Forcing a RIP update for a port or VLAN.....	127
Chapter 10: Common procedures using EDM.....	129
Configuring interVRF route redistribution policies.....	129

Chapter 11: Customer service.....	131
Getting technical documentation.....	131
Getting product training.....	131
Getting help from a distributor or reseller.....	131
Getting technical support from the Avaya Web site.....	132
Index.....	133

Chapter 1: New in this release

The following sections detail what's new in *Avaya Virtual Services Platform 9000 Configuration — OSPF and RIP*, NN46250–506 for Release 3.0:

- [Features](#) on page 7
- [Other changes](#) on page 7

Features

There are no feature changes.

Other changes

See the following sections for information about changes that are not feature-related.

Network loss of designated router

[OSPF with switch clustering](#) on page 46 is added to address a situation when the network loses the designated router in a Routed Split MultiLink Trunking configuration.

New in this release

Chapter 2: Introduction

Use this document to configure the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) on the Avaya Virtual Services Platform 9000. The router uses these protocols to determine the best routes for data forwarding.

For information about the Border Gateway Protocol, see *Avaya Virtual Services Platform Configuration — BGP Services, NN46250-507*.

- [Routing fundamentals](#) on page 11
- [OSPF fundamentals](#) on page 27
- [RIP fundamentals](#) on page 49
- [OSPF configuration using EDM](#) on page 53
- [OSPF configuration using ACLI](#) on page 83
- [RIP configuration using EDM](#) on page 107
- [RIP configuration using ACLI](#) on page 119
- [Common procedures using EDM](#) on page 129

Chapter 3: Routing fundamentals

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

For more information about how to use the Avaya command line interface (ACLI), see *Avaya Virtual Services Platform 9000 — User Interface Fundamentals, NN46250-103*.

- [Routing protocols](#) on page 11
- [IP addresses](#) on page 12
- [VLANs and routing](#) on page 16
- [Static routes](#) on page 17
- [Black hole static routes](#) on page 17
- [Circuitless IP](#) on page 18
- [Route policies](#) on page 19
- [IP routing features and considerations](#) on page 23

Routing protocols

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

The Avaya Virtual Services Platform 9000 supports wire-speed IP routing of frames using one of the following dynamic unicast IP routing protocols for path selection:

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

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

The VSP 9000 uses routing protocols like OSPF and RIP to populate routing tables. Routers use a routing protocol to exchange network topology information. A router uses the IP address of an incoming data packet to send the packet according to the routing tables.

The most commonly used unicast routing protocols include OSPF, RIP, and BGP. For more information about BGP, see *Avaya Virtual Services Platform 9000 Configuration — BGP Services, NN46250-507*. For information about multicast routing protocols, see *Avaya Virtual Services Platform 9000 Configuration — IP Multicast Routing Protocols, NN46250-504*.

IP addresses

An IP version 4 (IPv4) address consists of 32 bits expressed in dotted-decimal format (x.x.x.x). The IPv4 address space divides into classes, with classes A, B, and C reserved for unicast addresses. Class A, B, and C account for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. The following table describes IP address space by address range and mask.

Table 1: IP addresses

Class	Address range	Mask	Number of addresses
A	1.0.0.0—126.0.0.0	255.0.0.0	126
B	128.0.0.0—191.0.0.0	255.255.0.0	127 * 255
C	192.0.0.0—223.0.0.0	255.255.255.0	31 * 255 * 255
D	224.0.0.0—239.0.0.0	—	—

To express an IP address in dotted-decimal notation, convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary, has a different boundary point between the network and host portions of the address as illustrated in the following figure. The network portion is a network number from 8 to 24 bits. The remaining bits identify a specific host on the network.

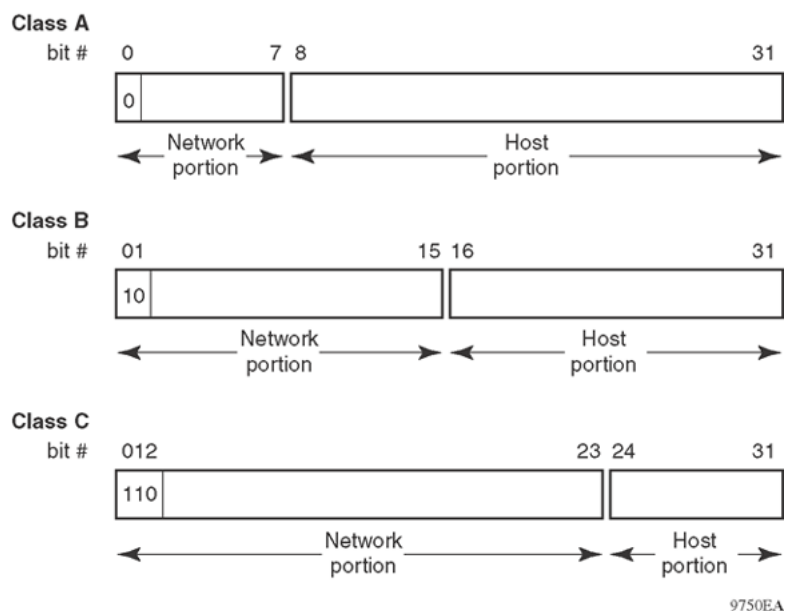


Figure 1: Network and host boundaries in IP address classes

IPv4 addresses are 32 bits long and expressed in decimal format.

Subnet addressing

Subnetworks (or subnets) extend the IP addressing scheme used by an organization to one with an IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

Create a subnet address by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is in the first octet of the host portion (10). A subnet mask applies to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks can create different numbers of subnets and hosts. This example includes using the zero subnet, which Virtual Services Platform 9000 supports.

Table 2: Subnet masks for class B and class C IP addresses

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
Class B			
2	255.255.192.0	2	16 382
3	255.255.224.0	6	8 190
4	255.255.240.0	14	4 094
5	255.255.248.0	30	2 046
6	255.255.252.0	62	1 022

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1 022	62
11	255.255.255.224	2 046	30
12	255.255.255.240	4 094	14
13	255.255.255.248	8 190	6
14	255.255.255.252	16 382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

With variable-length subnet masking (VLSM), you can divide your intranet into pieces that match your requirements. Routing is based on the longest subnet mask or network that matches. RIPv2 and OSPF both support VLSM.

Supernet addressing and CIDR

A supernet, or classless interdomain routing (CIDR) address, is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed. Using supernetting, you can address an entire block of class C addresses and avoid using large routing tables to track the addresses.

Each supernet has a unique supernet address that consists of the upper bits shared by all addresses in the contiguous block. For example, consider the class C addresses shown in the following figure. By adding the mask 255.255.128.0 to IP address 192.32.128.0, you aggregate the addresses 192.32.128.0 through 192.32.255.255, and 128 class C addresses use a single routing advertisement. As shown in the following figure, you use the address 192.32.0.0/17 to aggregate 128 addresses (192.32.0.0/24 to 192.32.127.0/24).




Figure 2: Class C address supernet

Another example is the block of addresses 192.32.0.0 to 192.32.7.0. The supernet address for this block is 11000000 00100000 00000000, with the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an address–mask pair:

- The address is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- The mask is a 32-bit string that contains a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

Although classes prohibit using an address mask with the IP address, you can use CIDR to create networks of various sizes using the address mask. You can also divide your address space using VLSM; however, the division is not used outside your network. With CIDR, routers outside your network use your addresses.

VLANs and routing

To route traffic on a virtual local area network (VLAN), you assign an IP address to the VLAN and not with a particular physical port. Brouter ports use single-port VLANs to route IP packets and bridge nonroutable traffic in specifically assigned VLANs.

Virtual routing between VLANs

The Avaya Virtual Services Platform 9000 supports wire-speed IP routing between VLANs. As shown in the following figure, although VLAN 1 and VLAN 2 are on the same switch, for traffic to flow from VLAN 1 to VLAN 2, the traffic must be routed.




Figure 3: IP routing between VLANs

To configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a virtual router interface address for the VLAN (a virtual router interface is so named because it is associated with no particular port). Through a VLAN port, you can reach the VLAN IP address, and the VLAN routes frames through the gateway IP address. The system forwards routed traffic to another VLAN within the switch.

If you use a spanning tree protocol on a VLAN, spanning tree convergence must be stable before the routing protocol begins. This requirement can lead to an additional delay in forwarding IP traffic.

Because a port can belong to multiple VLANs (some of which are routed on the switch and some of which are not), a one-to-one correspondence no longer exists between the physical port and the router interface.

As with an IP address, you can also use virtual router interface addresses for device management. For Simple Network Management Protocol (SNMP) or Telnet management, you can use a virtual router interface address to access the switch if you enable routing on the VLAN.

Brouter ports

The Virtual Services Platform 9000 supports 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 IP protocol-based VLAN that performs 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 and can still route IP traffic. This feature removes interruptions caused by spanning tree recalculations in routed traffic.

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

Static routes

You can use static routes to manually create routes to a destination IP address.

You can use a static default route to specify a route to all networks for which no explicit routes exist in the forwarding information base (FIB) or the routing table. This route is, by definition, a route with the prefix length of zero (RFC1812). You can configure the Virtual Services Platform 9000 with a route by using the IP static routing table.

To create a default static route, you must configure the destination address and subnet mask to 0.0.0.0.

You can configure a static route with a next-hop that does not directly connect, but that hop must be reachable. Otherwise, the static route is disabled.

Static routes and High Availability mode

High Availability (HA) mode supports the creation of static routes to enhance network stability. After you configure a static route on a master Control Processor (CP) module, the standby CP module uses the same configuration through synchronization. You can use the local next-hop option to configure a static route with or without a local next hop.

For more information about HA mode and support, see *Avaya Virtual Services Platform 9000 Administration, NN46250-600* and *Avaya Virtual Services Platform 9000 Release Notes — Software Release 3.0, NN46250-401*.

Black hole static routes

A black hole static route is a route with an invalid next hop. The switch drops packets destined for this network.

While it aggregates or injects routes to other routers, the router itself cannot have a path to the aggregated destination. In such cases, the result is a black hole or a routing loop. To avoid such loops, configure a black hole static route to the advertised destination.

You can configure a preference value for a black hole route. Configure that preference value so that the route is elected as the best route.

Before you add a black hole static route, ensure that no other static route to the same destination is enabled. If such a route exists, you cannot add the black hole route.

If you enable a black hole route, you cannot add another static route to that destination. You must delete or disable the black hole route before you add a regular static route to that destination.

Circuitless IP

Circuitless IP (CLIP) is a virtual (or loopback) interface that is not associated with a physical port. You can use the CLIP interface to provide uninterrupted connectivity to your switch if a path is available to reach the device.

A CLIP address, or a loopback address, is an IP address that is not tied to a specific interface. Because the CLIP address is not tied to a physical port or VLAN, the CLIP state is always active.

You can use a CLIP address as the OSPF router ID. If you use BGP with OSPF, the OSPF router ID becomes the BGP identifier automatically. Therefore, in this case, Avaya recommends that you use the CLIP address as the OSPF router ID. By doing so, the OSPF router ID is always active regardless of the port state (up or down).

For example, as shown in the following figure, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). An Interior Border Gateway Protocol (IBGP) session exists between two additional addresses, 195.39.128.1/30 (CLIP 1) and 195.39.281.2/30 (CLIP 2).




Figure 4: Routers with IBGP connections

CLIP 1 and CLIP 2 represent the virtual CLIP addresses between R1 and R2. These virtual interfaces are not associated with the physical link or hardware interface. The IBGP session can continue as long as a path exists between R1 and R2. An Interior Gateway Protocol (IGP), for example, OSPF, routes addresses that correspond to the CLIP addresses. After the routers in the autonomous system (AS) learn all the CLIP addresses, the routers establish the IBGP session and exchange routes.

The router treats the CLIP interface like an IP interface. The router treats the network associated with the CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

The router advertises routes to other routers in the domain either as external routes using the route redistribution process or after you enable OSPF in a passive mode to advertise an OSPF internal route. You can configure only the OSPF protocol on the circuitless IP interface.

After you create a CLIP interface, the system software programs a local route with the CPU as the destination ID. The CPU processes all packets destined to the CLIP interface address. The system treats other packets with destination addresses associated with this network (but not with the interface address) as if they are from an unknown host.

You can use a CLIP address as the source IP address in the IP header to send remote monitoring (RMON) traps.

Route policies

When the Virtual Services Platform 9000 routes IP traffic, you can apply a number of route policies (filters) to manage, accept, redistribute, and announce policies for unicast routing table information. The filtering process relies on the IP prefix lists in the common routing table manager infrastructure. Filters apply in different ways to various unicast routing protocols.

The following figure shows how filters apply to the BGP, RIP, and OSPF protocols.

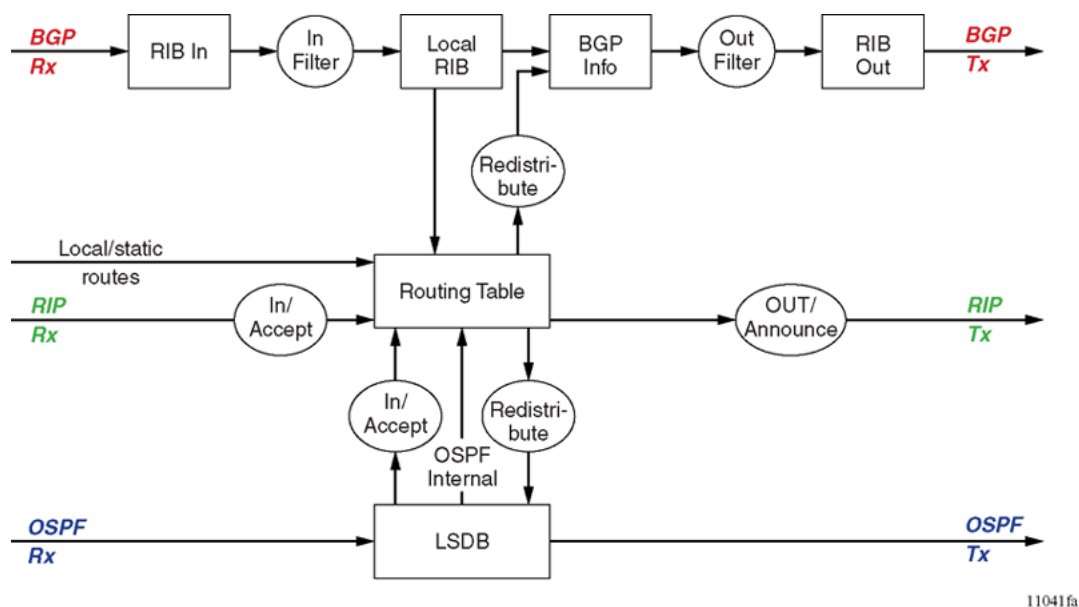


Figure 5: Route filtering for unicast routing protocols

Accept policies

Accept policies (or in filters) apply to incoming traffic to determine whether to add the route to the routing table. Accept policies apply in different ways to different protocols, as follows:

- RIP and BGP—filters apply to all incoming route information
- OSPF—filters apply only to external route information. Filters do not apply to internal routing information because other routers in the OSPF domain can have inconsistent databases that can affect the router view of the network topology.

In a network with multiple routing protocols, you can prefer specific routes from RIP instead of from OSPF. The network prefix is a commonly used match criterion for accept policies and ingress filters.

Use RIP accept policies to selectively accept routes from RIP updates. If you do not define policies, the default behavior applies, which adds all learned routes to the route table.

Use RIP accept policies to:

- Listen to RIP updates only from certain gateways.
- Listen only to specific networks.
- Assign a specific mask included with a network in the routing table (such as a network summary).

Redistribution and redistribution filters

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

Redistribution filters (policies) notify the routing protocol (within the device) of changes in the route table. In Virtual Services Platform 9000 software, announce policies performed interface-based redistribution. You must strictly apply announce policies to link-state advertisements (LSA), RIP updates, or BGP Network Layer Reachability Information (NLRI) to their respective domains. With redistribution filters, if you do not breach the protocol rules, you can choose not to advertise everything in the protocol database, or you can summarize or suppress route information. On the Virtual Services Platform 9000, by default, external routes do not leak to protocols you do not configure.

Announce policies

Announce policies (or out filters) apply to outgoing advertisements to neighbors and peers in the protocol domain to determine whether to announce specific route information. Out filtering applies to RIP updates and BGP NLRI updates.

In contrast, out filtering does not apply to OSPF information because OSPF routing information must always be consistent across the domain. To restrict the flow of external route information in the OSPF protocol database, apply redistribution filters instead of out filters.

Filter and policy application

The following figure shows the three distinct filtering stages that apply to IP routing protocol updates.


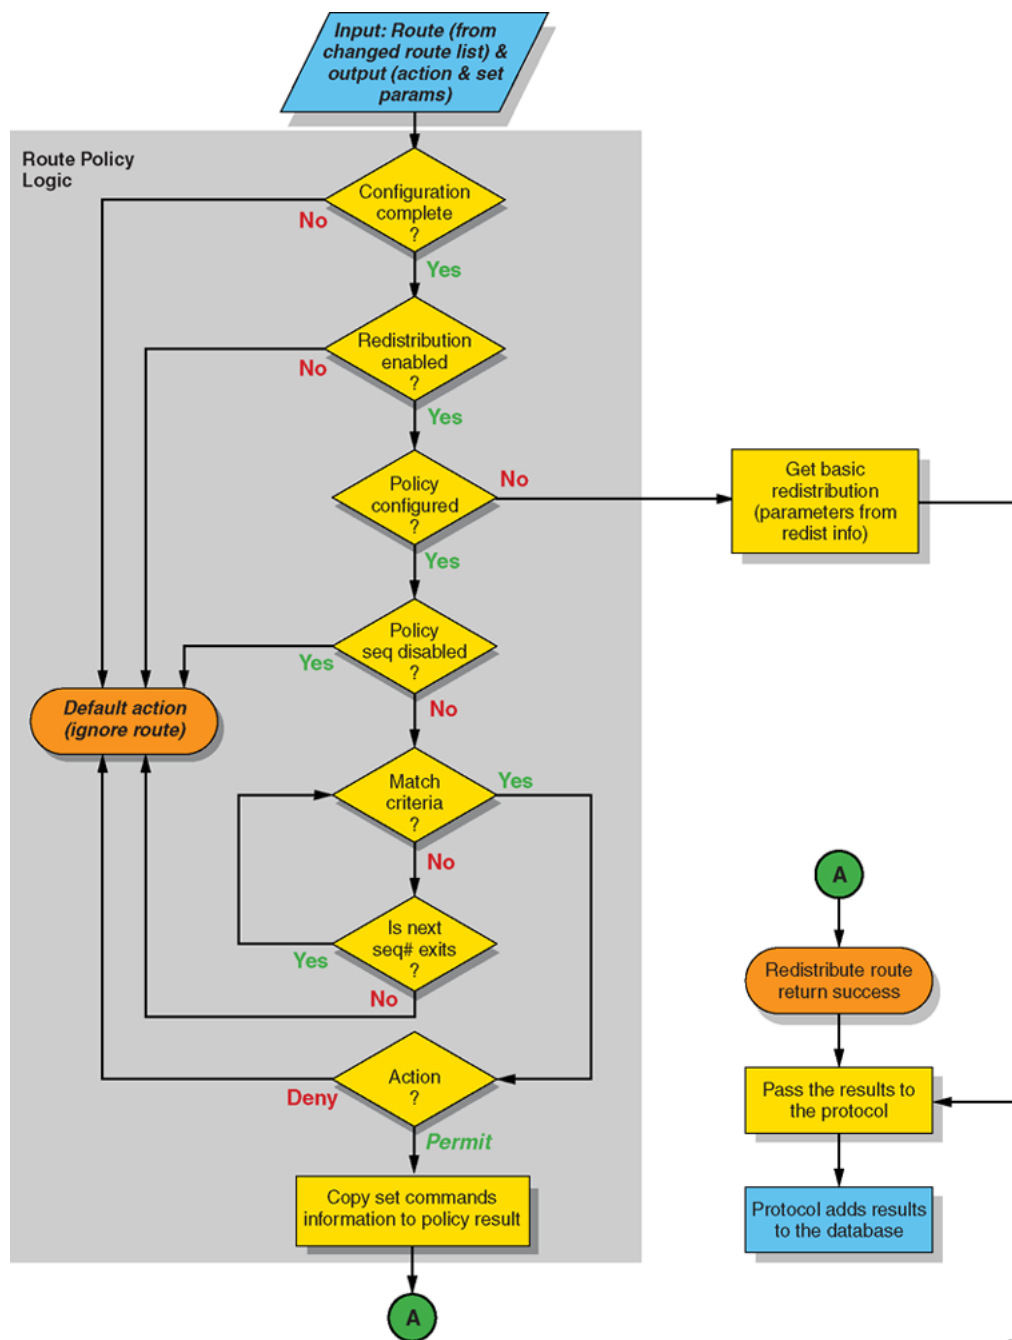


Figure 6: Route filtering stages

These stages are

1. Filter stage 1 is the accept policy or in filter. This filter applies to incoming routing protocol updates to detect changes in the dynamic (protocol-learned) routing information, which are submitted to the routing table.
2. Filter stage 2 is the redistribution filter. This filter applies to the entries in the protocol routing table during the route leak process.
3. Filter stage 3 is the announce policy or out filter. This filter applies to outgoing routing protocol updates within a protocol domain.

The following figure shows the logical operations that occur during the route-filtering process in the Virtual Services Platform 9000.



10533EB

Figure 7: Route filtering logic

For information about how to configure policies, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

IP routing features and considerations

The Virtual Services Platform 9000 provides features that you can use to maximize routing efficiency. This section contains information that you can use to help you configure IP routing.

Equal Cost Multipath

You can use multiple paths for load sharing of traffic. These multiple paths provide fast convergence to other active paths if the network fails. By maximizing load sharing among equal-cost paths, you can use your links between routers to efficiently send IP traffic. Equal Cost Multipath (ECMP) supports and complements the following protocols and route types:

- OSPF
- RIP
- BGP
- static routes
- default routes

For more information about ECMP, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

Alternative routes and route preferences

Routers can learn several routes to a destination network through several protocols. In the Virtual Services Platform 9000 software, if you enable the alternative route function, the switch stores the alternative routes, sorted in order of network mask, cost, and route preference. The first route on this list is the best route. The hardware uses the best route. Other routes are alternative routes.

To avoid traffic interruption, you can globally enable alternative routes to replace the best route with the next-best route if the best route becomes unavailable. Alternative routes apply between routing protocols. For example, if an OSPF route becomes unavailable and an alternate RIP route is available, the router immediately activates the alternate route without waiting for an update interval to expire.

The internal routing table manager (RTM) records the route changes for protocols. The RTM maintains separate tables of static (user-configured) and dynamic (protocol-learned) routes. You can configure preferences that determine the precedence assigned to one type of route over another.

If a router learns a route with the same network mask and cost values from multiple sources (protocols), the router uses preferences to select the best route. The router holds up to four other routes for each destination as alternative routes.

You can configure route preferences for static routes and routing protocols. When you configure a static route, you can specify a preference for the route. To modify the preference for a static route, disable the route before you edit the configuration, and then reenabling the route.



Important:

Changing route preferences is a process-intensive operation that can affect system performance and network reachability while you perform route preference procedures. Avaya recommends that if you want to change preferences for static routes or routing protocols, do so when you configure routes or during a maintenance window.

All standard routing protocols use a default preferences. You can modify the default preference for a protocol to give it higher or lower priority than other protocols. After you change the preference for a route, if all best routes remain best routes, only the local route tables change. However, if a change in the protocol preference causes best routes to no longer be best routes, this change can affect neighboring route tables.

In addition, you can modify the preference value for dynamic routes through route filters or IP policies, and this value overrides the global preference for the protocol. You can use alternative mechanisms to change the behavior of specific routes to have a different preference rather than by acquiring the global protocol preference. For a static route, you can specify an individual route preference that overrides the global static route preference. The preference value can be a number from 0 to 255, with 0 reserved for local routes. 255 represents an unreachable route.

Reverse path checking

Reverse path checking prevents packet forwarding for incoming IP packets that have incorrect or forged (spoofed) IP addresses. Reverse path checking guarantees that traffic received on one interface was sent by a station from a specific interface (which prevents address spoofing). With this mode enabled, the Virtual Services Platform 9000 performs a reverse path check to verify the packet source IP address. If the switch cannot verify the source, the switch discards the packet.

Configure reverse path checking as required for each IP interface. When enabled, the Virtual Services Platform 9000 checks all routing packets that enter the interface. Reverse path checking ensures that the source address and source interface appear in the routing table, and that the address matches the interface that receives the packet.

You can use one of two modes for reverse path checking:

- **Exist-only mode:** Reverse path checking checks whether the source IP address for the incoming packet exists in the routing table. If the switch finds the source IP entry, the switch forwards the packet as usual; otherwise, the switch discards the packet.
- **Strict mode:** Reverse path checking checks that the source IP address exists in the routing table, and is reachable through the incoming IP interface (and not through other interfaces). If these conditions are not met, the switch discards the packet.

For more information about how to configure Reverse path checking, see *Avaya Virtual Services Platform 9000 Security, NN46250-601*.

The following example illustrates how strict mode works.




Figure 8: Reverse path checking network configuration

Consider the following parameters:

- A router connects a server (32.57.5.10) to a client (192.32.45.10).
- The router uses reverse path checking.
- The router has the following entries in the routing table:

Table 3: Routing table

Destination address	Next-hop address	Forward through port
32.57.5.10	173.56.42.2	3/7
192.32.45.10	145.34.87.2	7/2
192.32.46.10	145.34.88.2	7/1

If the client sends a legitimate packet, the following actions occur:

- The client sends packet to the server. The packet has a source IP address of 192.32.45.10 and a destination IP address of 32.57.5.10.
- The packet arrives at router port 7/2 (brouter). The routing engine performs a destination IP address lookup and finds the destination port is 3/7.
- Reverse path checking begins. The routing engine searches for the source IP address of 192.32.45.10. The routing engine finds an entry in the routing table that specifies the next-hop port as 7/2, which matches the packet incoming port. Because the address and port information matches, the switch forwards the packet as usual.

If the client sends a spoofed packet, the following actions occur:

- The client sends a packet to the server with a forged IP address of 192.32.46.10 through port 7/2.
- Reverse path checking finds that the source IP address next-hop port is 7/1, which does not match the packet incoming port of 7/2. In this case, the switch discards the packet.

You can think about reverse path checking as follows. If A sends packets to B through route X ingress port Y, then the return packets from B to A must egress X through the same port Y. If returning packets take a different path, the switch drops them.

Chapter 4: OSPF fundamentals

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

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

For information about the Border Gateway Protocol (BGP), see *Avaya Virtual Services Platform 9000 Configuration — BGP Services, NN46250-507*

- [OSPF overview](#) on page 27
- [Dijkstras algorithm](#) on page 28
- [Autonomous system and areas](#) on page 29
- [OSPF neighbors](#) on page 32
- [Router types](#) on page 33
- [OSPF interfaces](#) on page 34
- [OSPF and IP](#) on page 39
- [OSPF packets](#) on page 39
- [Intra-area link-state advertisements](#) on page 40
- [ASE routes](#) on page 41
- [OSPF virtual links](#) on page 41
- [OSPF ASBRs](#) on page 42
- [OSPF metrics](#) on page 44
- [OSPF security mechanisms](#) on page 44
- [OSPF and route redistribution](#) on page 45
- [OSPF configuration considerations](#) on page 46
- [OSPF with switch clustering](#) on page 46

OSPF overview

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

AS. Routers synchronize their topological databases based on the sharing of information from LSAs.

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

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

In large networks, OSPF offers the following benefits:

- fast convergence

After network topology changes, OSPF recalculates routes quickly.

- minimal routing protocol traffic

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

- load sharing

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

- type of service

OSPF can calculate separate routes for each IP TOS.

Dijkstras algorithm

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

1. After a router starts, it initializes the OSPF data structures, and then waits for indications from lower-level protocols that the router interfaces are functional.
2. A router then uses the Hello protocol to discover neighbors. On point-to-point and broadcast networks the router dynamically detects neighbors by sending hello packets to the multicast address AllSPFRouters. On Non-Broadcast Multiple Access (NBMA) networks, you must provide some configuration information to discover neighbors.
3. On all multiaccess networks (broadcast or nonbroadcast), the Hello protocol elects a designated router (DR) for the network.
4. The router attempts to form adjacencies with some of its neighbors. On multiaccess networks, the DR determines which routers become adjacent. This behavior does not occur if you configure a router as a passive interface because passive interfaces do not form adjacencies.

5. Adjacent neighbors synchronize their topological databases.
6. The router periodically advertises its link state, and does so after its local state changes. LSAs include information about adjacencies, enabling quick detection of dead routers on the network.
7. LSAs flood throughout the area to ensure that all routers in an area have an identical topological database.
8. From this database each router calculates a shortest-path tree, with itself as the root. This shortest-path tree in turn yields a routing table for the protocol.

Autonomous system and areas

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

You can attach a router to more than one area. When you perform this action, you can maintain a separate topological database for each connected area. Two routers within the same area maintain an identical topological database for that area. Each area uses a unique area ID and the area ID 0.0.0.0 is reserved for the backbone area.

The router routes packets in the AS based on their source and destination addresses. If the source and destination of a packet reside in the same area, the router uses intra-area routing. If the source and destination of a packet reside in different areas, the router uses inter-area routing. Intra-area routing protects the area from bad routing information because it does not use routing information obtained from outside the area. Inter-area routing must pass through the backbone area. For more information about the backbone area, see [Backbone area](#) on page 30.

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

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

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

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

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

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

Backbone area

The backbone area consists of the following network types:

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

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

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

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

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

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

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

Stub area

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

Not so stubby area

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

In an OSPF NSSA, the NSSA N/P bit notifies the ABR which external routes to advertise to other areas. If the NSSA N/P bit is set (the value is 1), the ABR exports the external route. This configuration is the default for the Avaya Virtual Services Platform 9000. When the NSSA N/

P bit is not set (the value is 0), the ABR drops the external route. You can create a route policy on the Virtual Services Platform 9000 to manipulate the N/P bit.

Multiarea OSPF configuration

The following figure shows five Virtual Services Platform 9000 devices (R1 to R5) in a multi-area configuration.

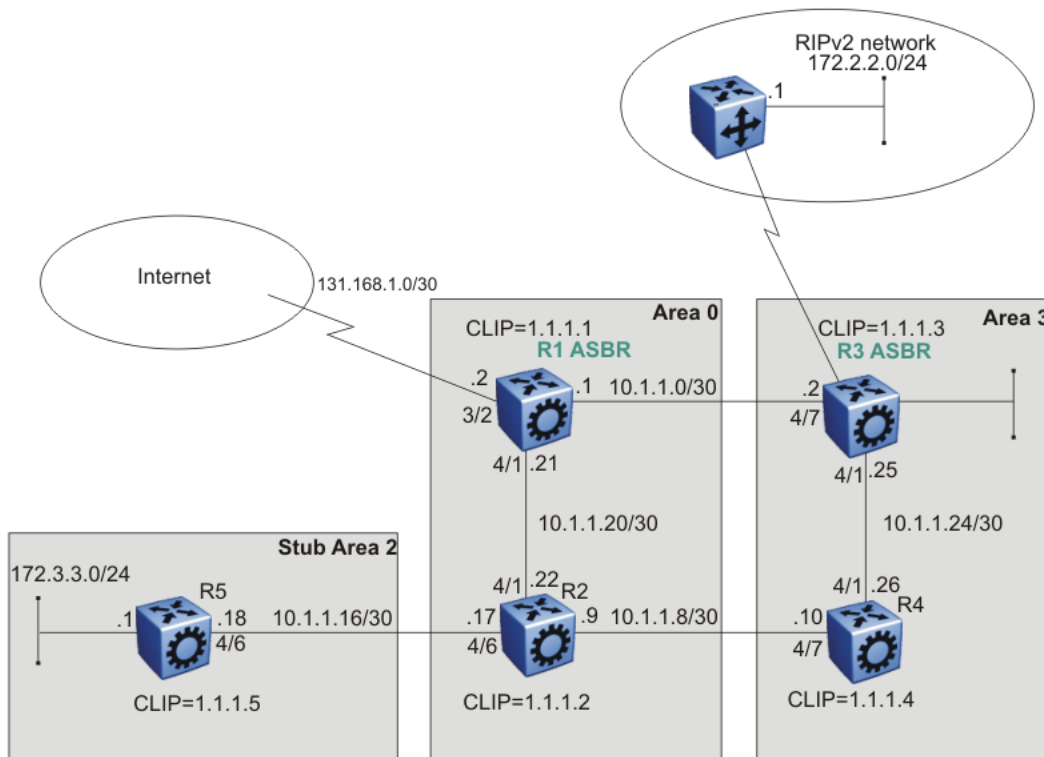


Figure 9: Multiarea configuration example

The following list explains the configuration for the Virtual Services Platform 9000 devices R1 through R5:

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

Network 172.3.3.0/24 on R5 uses a VLAN configuration instead of a brouter port. This example uses brouter ports rather than VLANs because the spanning tree algorithm is disabled by default if you use brouter interfaces.

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

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

OSPF neighbors

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

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

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

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

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

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

Neighbors can form an adjacency to exchange routing information. After two routers form an adjacency, they perform a database exchange process to synchronize their topological databases. After the databases synchronize, the routers are fully adjacent. Adjacency

conserves bandwidth because, from this point, the adjacent routers pass only routing change information.

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

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

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

Router types

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

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

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

Table 4: Router types in an OSPF network

Router type	Description
AS boundary router	A router that attaches at the edge of an OSPF network is an ASBR. An ASBR generally has one or more interfaces that run an interdomain routing protocol such as Border Gateway Protocol. In addition, a router that distributes static routes or RIP routes into OSPF is an ASBR. The ASBR forwards external routes into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area border router	A router that attaches to two or more areas inside an OSPF network is an ABR. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information.

Router type	Description
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is an IR. Unlike ABRs, IRs have topological information only about the area in which they reside.
Designated router	In a broadcast or NBMA network, the routers elect a single router as the DR for that network. A DR makes sure that all routers on the network synchronize and advertises the network to the rest of the AS.
Backup designated router	A BDR is elected in addition to the DR and, if the DR fails, can assume the DR role quickly.

OSPF interfaces

Configure an OSPF interface, or link, on an IP interface. In the Virtual Services Platform 9000, an IP interface can be either a single link (brouter port) or a logical interface configured on a VLAN (multiple ports). The state information associated with the interface is obtained from the underlying lower-level protocols and the routing protocol itself.



Important:

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

OSPF network types allow OSPF-neighboring between routers over various types of network infrastructures. You can configure each interface to support various network types. The following table describes the OSPF network interface types supported by the Virtual Services Platform 9000.

Table 5: OSPF network types

Network interface type	Description
Broadcast interfaces on page 35	Broadcast interfaces automatically discover every OSPF router on the network by sending OSPF hello packets to the multicast group AllSPFRouters (224.0.0.5). Neighboring is automatic and requires no configuration.
Non-Broadcast Multiple Access interfaces on page 35	The NBMA network type models network environments that do not have native Layer 2 broadcast or multicast capabilities, such as Frame Relay and X.25. OSPF hello packets are unicast to manually configured neighbors.

Network interface type	Description
Passive interfaces on page 39	A passive interface is an interfacing network in OSPF that does not generate LSAs or form adjacencies. Use a passive interface on an access network or on an interface used for BGP peering. Using passive interfaces limits the amount of CPU cycles required to perform the OSPF routing algorithm.

Broadcast interfaces

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

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

Non-Broadcast Multiple Access interfaces

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

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

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

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

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




Figure 10: NBMA subnet

NBMA interface operations and parameters

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

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

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

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

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

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

RouterDeadInterval period, the router sends hello packets at the specified poll interval, for example, every 120 seconds.

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

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

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

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

OSPF and NBMA example: adjacency formation

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

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

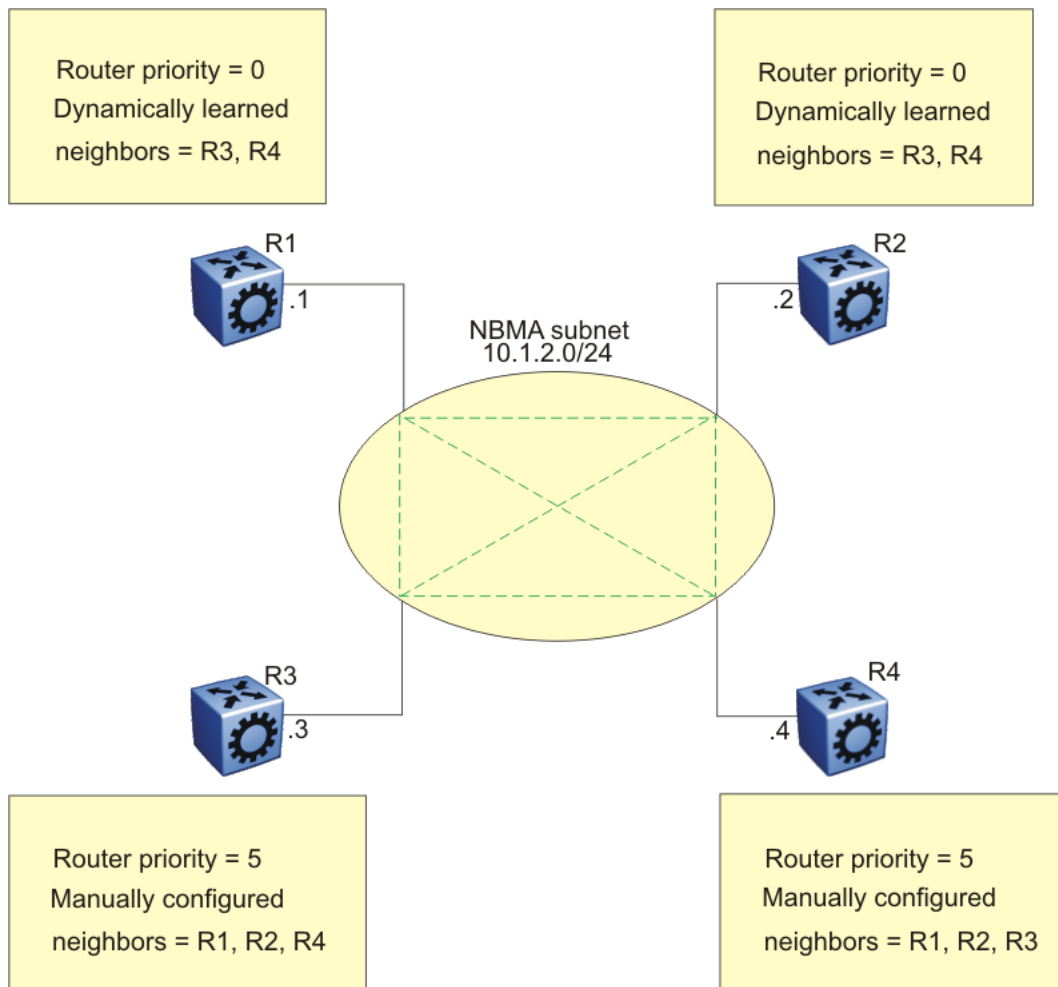


Figure 11: NBMA subnet configuration example

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

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

To create this NBMA network, configure the following parameters:

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

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

1. R3 and R4 send each other a hello packet to elect a DR.
2. The Hello protocol elects R3 as the DR, and R4 as the BDR.

3. R3 (DR) and R4 (BDR) send hello packets to all other routers on the NBMA subnet to synchronize their link-state databases and establish themselves as DR and BDR.
4. R1 and R2 reply to R3 and R4.
5. R3 and R4 each form three adjacencies (one with each router on the NBMA subnet).
6. R1 and R2 each form two adjacencies (one with the DR and one with the BDR).

Passive interfaces

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

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

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

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

OSPF and IP

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

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

Because OSPF can specify a range of networks, it can send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 255.255.0.0 describes a single route to destinations 128.185.0.0 to 128.185.255.255.

OSPF packets

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

- The router transmits hello packets between neighbors and never forwards them. The Hello protocol requires routers to send hello packets to neighbors at predefined hello intervals. A neighbor router that does not receive a hello packet declares the other router dead.
- The router exchanges DD packets after neighboring routers establish a link, which synchronizes their LSDBs.
- Link-state request packets describe one or more link-state advertisements that a router requests from its neighbor. Routers send link-state requests if the information received in DD packets from a neighbor is not consistent with its own link-state database.
- Link-state update packets contain one or more LSAs and the router sends them following a change in network conditions.
- The router sends link-state acknowledgement packets to acknowledge receipt of link-state updates. Link-state acknowledgement packets contain the headers of the received LSAs.

Intra-area link-state advertisements

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

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

ASE routes

OSPF considers the following routes as ASE routes:

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

OSPF virtual links

On an OSPF network, a Virtual Services Platform 9000 that acts as an ABR must connect directly to the backbone. If no physical connection is available, you can automatically or manually configure a virtual link.

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

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

On the Virtual Services Platform 9000, OSPF behavior follows OSPF standards; the router cannot learn OSPF routes through an ABR unless the ABR connects to the backbone or through a virtual link.

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




Figure 12: Virtual link between ABRs through a transit area

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

OSPF ASBRs

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

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

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

Table 6: ASBR routing hierarchy

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

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

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

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

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

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

Area link-state advertisements

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

Table 7: OSPF LSA types

LSA type	Description	Area of distribution
1	A router originates type 1 LSAs (router LSAs) to describe its set of active interfaces and neighbors.	Passed only within the same area
2	Type 2 LSAs (network LSAs) describe a network segment such as broadcast or NBMA. In a broadcast network, the DR originates network LSAs.	Passed only within the same area
3	The ABR originates type 3 LSAs (network-summary LSAs) to describe the networks within an area.	Passed between areas
4	Type 4 LSAs (ASBR-summary LSAs) advertise the location of the ASBRs from area to area.	Passed between areas
5	Type 5 LSAs (ASE LSAs) describe networks outside of the OSPF domain. The ASBR originates type 5 LSAs. In stub and NSSA areas, a single default route replaces type 5 LSA routes.	Passed between areas

LSA type	Description	Area of distribution
6	Type 6 LSAs (group membership LSAs) identify the location of multicast group members in multicast OSPF.	Passed between areas
7	Type 7 LSAs import external routes in OSPF NSSAs.	Translated between areas

OSPF metrics

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

Assign default metric speeds for different port types, such as 10 Mb/s or 1 Mb/s ports. On Virtual Services Platform 9000, you can specify a new metric speed for an IP interface. An IP interface can be a brouter port or a VLAN.

RFC1583 states the following:

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

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

OSPF security mechanisms

The Virtual Services Platform 9000 implementation of OSPF includes security mechanisms to prevent unauthorized routers from attacking the OSPF routing domain. These security mechanisms prevent a malicious person from joining an OSPF domain and advertising false information in the OSPF LSAs. Likewise, security prevents a misconfigured router from joining an OSPF domain.

Simple password

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

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

Message Digest 5

Avaya recommends that you use Message Digest 5 (MD5) for OSPF security because it provides standards-based (RFC1321) authentication using 128-bit encryption. When you use MD5 for OSPF security, it is almost impossible for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

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

OSPF and route redistribution

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

You can redistribute routes

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

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

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

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

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

OSPF configuration considerations

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

OSPF host route advertisements and nonbackbone areas

The Virtual Services Platform 9000 does not associate a host route with a specific area. Therefore, if you create a host route in a nonbackbone area, nonbackbone (nonOSPF core) areas do not advertise it.

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

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

OSPF with switch clustering

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

Each router performs the SPF run as soon as it detects a new DR. Depending on the speed of the router, the router can perform the SPF run before it receives the new LSAs for the segments, which requires a second SPF run to update and continue routing across the segment. The OSPF hold-down timer does not permit 2 consecutive SPF runs within the value of the timer. This limitation can lead to traffic interruption of up to 10 seconds.

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

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

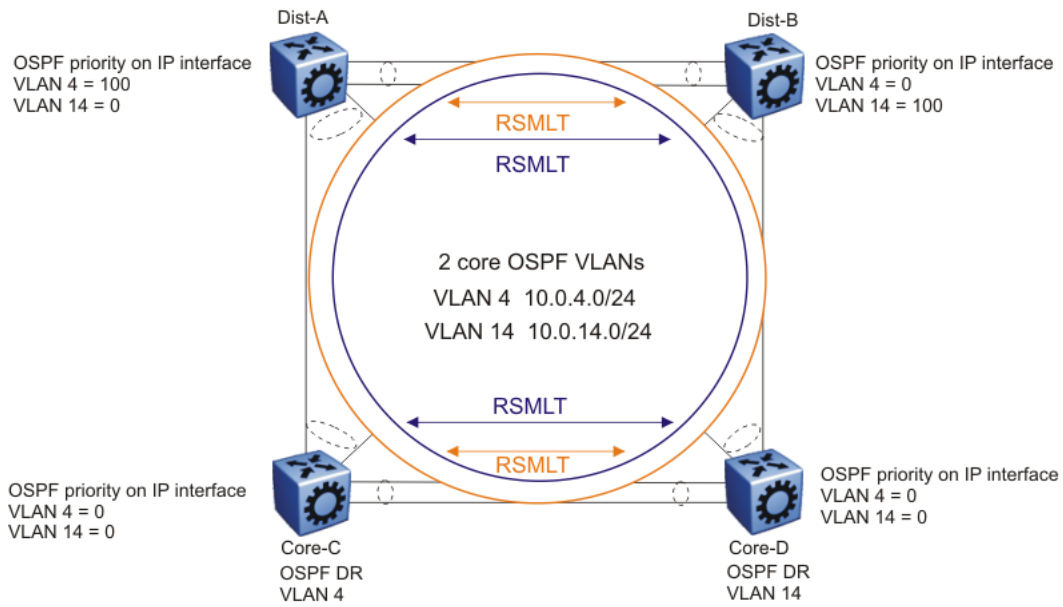


Figure 13: RSMLT with dual core VLANs

Chapter 5: RIP fundamentals

Use the information in these sections to help you understand the Routing Information Protocol (RIP) .

For more information about the Border Gateway Protocol (BGP), see *Avaya Virtual Services Platform 9000 Configuration — BGP Services, NN46250-507*

- [Routing Information Protocol](#) on page 49
- [RIP and route redistribution](#) on page 50

Routing Information Protocol

In routed environments, routers communicate with one another to track available routes. Routers can dynamically learn about available routes using the RIP. The Avaya Virtual Services Platform 9000 software implements standard RIP to exchange IP route information with other routers.

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

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




Figure 14: Hop count or metric in RIP

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

The Virtual Services Platform 9000 supports RIPv2, which advertises routing table updates using multicast instead of broadcasting. RIPv2 supports variable length subnet masks (VLSM) and triggered router updates. RIPv2 sends mask information. If RIP does not receive information about a network for 90 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 180 seconds (six update intervals), it removes the network from the routing table. You can change the default timers by configuring the RIP interface timeout timer and the holddown timer.

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

RIP and route redistribution

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

You can redistribute routes

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

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

You can redistribute routes on a global basis, rather than for every interface. Virtual Services Platform 9000 adds support for global RIP redistribution. Use the `ip rip redistribute` command to

accomplish the (intraVRF) redistribution of routes through RIP, so that RIP redistribution occurs globally on all RIP-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

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

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

Chapter 6: OSPF configuration using EDM

Configure Open Shortest Path First (OSPF) parameters so that the switch can participate in OSPF routing operations. The following section describes procedures that you use while you configure OSPF on the Avaya Virtual Services Platform 9000 using Enterprise Device Manager (EDM).

- [Configuring OSPF globally](#) on page 54
- [Enabling OSPF globally](#) on page 56
- [Configuring global default metrics](#) on page 56
- [Configuring an OSPF interface](#) on page 57
- [Changing an OSPF interface type](#) on page 60
- [Configuring NBMA interface neighbors](#) on page 61
- [Configuring OSPF interface metrics](#) on page 62
- [Viewing all OSPF-enabled interfaces](#) on page 63
- [Configuring OSPF on a port](#) on page 64
- [Configuring OSPF on a VLAN](#) on page 66
- [Creating stubby or not-so-stubby OSPF areas](#) on page 69
- [Configuring stub area metrics advertised by an ABR](#) on page 70
- [Inserting OSPF area aggregate ranges](#) on page 71
- [Enabling automatic virtual links](#) on page 73
- [Configuring a manual virtual interface](#) on page 73
- [Viewing virtual neighbors](#) on page 75
- [Configuring host routes](#) on page 76
- [Enabling ASBR status](#) on page 77
- [Managing OSPF neighbors](#) on page 78
- [Viewing the link-state database](#) on page 78
- [Viewing the external link-state database](#) on page 79
- [Configuring route redistribution to OSPF](#) on page 80
- [Forcing shortest-path calculation updates](#) on page 82

Configuring OSPF globally

Prerequisites

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.
- Assign an IP address to the switch.

Configure OSPF parameters, such as automatic virtual links and OSPF metrics, to control how OSPF behaves on the system.

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

Table 8: Variable definitions

Variable	Value
RouterId	Specifies the OSPF router ID. This variable has the same format as an IP address but distinguishes this router from other routers in the OSPF domain.
AdminStat	Shows the administrative status of OSPF for the router. Enabled denotes that the OSPF process is active on at least one interface; disabled disables it for all interfaces. The default is disabled.
VersionNumber	Specifies the OSPF version.
AreaBdrRtrStatus	Denotes if this router is an area border router (ABR). AreaBdrRtrStatus value must be true to create a virtual router interface.
ASBdrRtrStatus	Specifies ASBR status. If you select the ASBdrRtrStatus check box, the router is an autonomous system boundary router (ASBR).
ExternLsaCount	Shows the number of external (LS type 5) link-state advertisements in the link-state database.
ExternLsaCksumSum	Shows the 32-bit unsigned sum of the link-state checksums of the external link-state advertisements in the link-state database. This sum determines if a change occurred in a router link-state database and compares the link-state databases of two routers.
OriginateNewLsas	Shows the number of new link-state advertisements originated from this router. This number increments each

Variable	Value
	time the router originates a new link-state advertisement (LSA).
RxNewLsas	Shows the number of received link-state advertisements that are new instances. This number does not include new instances of self-originated link-state advertisements.
10MbpsPortDefaultMetric	Indicates the default cost applied to 10 Mb/s interfaces (ports). The default is 100.
100MbpsPortDefaultMetric	Indicates the default cost applied to 100 Mb/s interfaces (ports). The default is 10.
1000MbpsPortDefaultMetric	Indicates the default cost applied to 1000 Mb/s interfaces (ports). The default is 1.
10000MbpsPortDefaultMetric	Indicates the default cost applied to 10 000 Mb/s interfaces (ports). The default is 1.
TrapEnable	Indicates whether to enable traps for OSPF. The default is false.
AutoVirtLinkEnable	Enables or disables the automatic creation of virtual links. The default is false.
SpfHoldDownTime	Specifies the OSPF holddown timer (3–60 seconds). The default is 10 seconds. The holddown timer delays a metric change due to a routing table update by x seconds. If you configure the timer to 0, OSPF accepts a new metric change immediately.
OspfAction	Initiates a new Shortest Path First (SPF) run to update the routing table. The default is none.
Rfc1583Compatability	Controls the preference rules used when the router chooses among multiple autonomous system external (ASE) LSAs which advertise the same destination. If enabled, the preference rule is the same as that specified by RFC1583. If disabled, the preference rule is as described in RFC2328, which can prevent routing loops when ASE LSAs for the same destination originate from different areas. The default is disable.
LastSpfRun	Indicates the time since the last SPF calculation made by OSPF.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **General** tab.
 4. Specify the OSPF router ID.

5. In AdminStart click the **enabled** option button.
 6. If required, configure the metrics that OSPF uses for 10, 100, 1000, and 10 000 Mb/s links.
The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
 7. To enable the switch to use OSPF SNMP traps, select the **TrapEnable** check box.
 8. To enable the automatic creation of virtual links, select the **AutoVirtLinkEnable** check box.
 9. Configure the OSPF holddown timer as required.
 10. Click **Apply**.
-

Enabling OSPF globally

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Configuring global default metrics

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs.

For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure the metrics that OSPF uses for 10, 100, 1000, and 10 000 Mb/s links. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.

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

Configuring an OSPF interface


Prerequisites

- Enable OSPF globally.
- Ensure that the interface exists (the port or VLAN has an IP address).
- You must know the network OSPF password to use password authentication, .
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure OSPF parameters, such as authentication and priority, to control how OSPF behaves on the interface. You can specify the interface as passive, broadcast, or Non-Broadcast Multiple Access (NBMA).

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

Table 9: Variable definitions

Variable	Value
IP Address	Specifies the IP address of the current OSPF interface
AddressLessIf	Designates whether an interface has an IP address: Interfaces with an IP address = 0 Interfaces without IP address = ifIndex
AreaId	Specifies the OSPF area name in dotted-decimal format. For VLANs, keeping the default area setting on the interface causes link-state database (LSDB) inconsistencies. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdminStat	Specifies the current administrative state of the OSPF interface (enabled or disabled).
RtrPriority	Specifies the OSPF priority to use during the election process for the designated router. The interface with the highest priority becomes the designated router. The interface with the second-highest priority becomes the backup designated router. If the priority is 0, the interface cannot become the designated router or the backup. The range is 0–255. The default is 1.
Type	Specifies the type of OSPF interface (broadcast or NBMA).  Important: To make it passive, first create the interface. After interface creation, click VLAN, VLANs to select the VLAN that is created with the OSPF interface. Click the IP tab and select the IP interface that is created with the OSPF interface. Lastly, click the OSPF tab and select Passive for the IfType .
AuthType	Specifies the type of authentication required for the interface. <ul style="list-style-type: none"> • none—No authentication required. • simple password—All OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—All OSPF updates received by the interface must contain the MD5 key. Configure the MD5 key in ACLI.
AuthKey	Specifies the key (up to 8 characters) required when you specify simple password authentication in the AuthType parameter.

Variable	Value
HelloInterval	Specifies the length of time, in seconds, between hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
TransitDelay	Specifies the length of time, in seconds, required to transmit an LSA update packet over the interface. The default is 1.
RetransInterval	Specifies the length of time, in seconds, required between LSA retransmissions. The default is 5.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the Hello interval. To avoid interpretability issues, the RtrDeadInterval value for the OSPF interface must match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. The default is 120.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Interfaces** tab.
 4. Click **Insert**.
 5. Select the IP address for the interface from the IP Address list.
 6. To designate a router priority, in the **RtrPriority** box, type a new value.
 7. In the **Type** area, select the type of OSPF interface you want to create.
 8. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.
 9. To enable authentication, in the **AuthType** area, select either **simplePassword** or **MD5**.
 10. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
 11. Click **Insert**.
 12. On the **Interfaces** tab, click **Apply**.
-

Changing an OSPF interface type

Prerequisites

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- If the interface is currently an NBMA interface with manually configured neighbors, you must first delete all manually configured neighbors.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Important:

The procedure above details the creation of a non-passive interface. Perform the following steps to create a passive interface:

- a. In the navigation tree, select **Configuration, VLAN, VLANs**.
- b. Click on the VLAN where the OSPF interface is created.
- c. Click the **IP** tab.
- d. Click on the IP Address where the OSPF interface is created.
- e. Click the **OSPF** tab.
- f. Click on Enable to disable the OSPF interface that was already enabled.

- g. Click **Apply**.
- h. Modify the interface type to passive.
- i. Click **Apply**.

Configuring NBMA interface neighbors

Prerequisites

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Ensure that the interface type is NBMA.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 10: Variable definitions

Variable	Value
IP Address	Specifies the neighbor IP address.
AddressLessIndex	Indicates addressed and addressless interfaces. This value is 0 on an interface with an IP address. On addressless interfaces, the corresponding value of ifIndex in the Internet standard management information base (MIB).
Router	Specifies the router ID of the neighboring router. The router ID has the same format as an IP address but identifies the router independent of its IP address.
Options	Specifies the bit mask that corresponds to the neighbor options parameter.
Priority	Specifies the priority.
State	Specifies the OSPF interface state.
Events	Specifies the number of state changes or error events that occur between the OSPF router and the neighbor router.

Variable	Value
Retransmission Queue Length	Specifies the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
ospfNbmaNbrPermanence	Indicates whether the neighbor is a manually configured NBMA neighbor; permanent indicates it is an NBMA neighbor.
HelloSuppressed	Indicates whether hello packets to a neighbor are suppressed.

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

Configuring OSPF interface metrics

Prerequisites

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 11: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **If Metrics** tab.
 4. Double-click the value cell, and type a new value.
 5. Click **Apply**.
When you enable a port for OSPF routing, the default metric in the port tab is 0. A value of 0 means that the port uses the default metrics for port types that you specify on the OSPF General tab.
-

Viewing all OSPF-enabled interfaces

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

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

Configuring OSPF on a port

Prerequisites

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

Configure OSPF parameters on a port to control how OSPF behaves on the port.

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

Table 12: Variable definitions

Variable	Value
Enable	Enables or disables OSPF routing on the specified port. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet, and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this port in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for the type of service (TOS) on this port. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1.

Variable	Value
	<ul style="list-style-type: none"> • FFFF—No route exists for this TOS. • IPCP links—Defaults to 0. • 0—Use the interface speed as the metric value when the state of the interface is up.
AuthType	<p>Specifies the type of authentication required for the port.</p> <ul style="list-style-type: none"> • none—No authentication required. • simple password—All OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—All OSPF updates received by the interface must contain the MD5 key. Configure the MD5 key in ACLI.
AuthKey	<p>Specifies the key (up to 8 characters) when you specify simple password authentication in the port AuthType variable.</p>
AreaId	<p>Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).</p>
AdvertiseWhenDown	<p>Advertises the network on this port as up, even if the port is down. The default is false.</p> <p>After you configure a port with no link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even if the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.</p>
IfType	<p>Specifies the type of OSPF interface (broadcast, NBMA, or passive).</p> <p>Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.</p>
PollInterval	<p>Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must have the same poll interval.</p>
IfMtuIgnore	<p>Specifies whether the interface ignores the global maximum transmission unit (MTU) configuration. To allow the Virtual Services Platform 9000 to accept OSPF database description (DD) packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.</p>

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

Configuring OSPF on a VLAN

Prerequisites

- Enable OSPF globally.
- Ensure that the VLAN uses an IP address.
- Ensure that the `ospf_md5key.txt` file is on the switch to use MD5 authentication.
- Ensure that you know the network OSPF to use password authentication, .
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure OSPF parameters on a VLAN to control how OSPF behaves on the VLAN.

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

Table 13: Variable definitions

Variable	Value
Enable	Enables or disables OSPF routing on the specified VLAN. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this VLAN in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for this TOS on this VLAN. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. <ul style="list-style-type: none"> • FFFF—No route exists for this TOS. • IPCP links—Defaults to 0. • 0—Use the interface speed as the metric value when the state of the interface is up.
AuthType	Specifies the type of authentication required for the VLAN. <ul style="list-style-type: none"> • none—No authentication required. • simple password—All OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—All OSPF updates received by the interface must contain the MD5 key.

Variable	Value
AuthKey	Specifies the key (up to eight characters) when you specify simple password authentication in the VLAN AuthType variable.
AreaId	Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	Advertises the network even if the port is down. If true, OSPF advertises the network on this VLAN as up, even if the port is down. The default is false. After you configure a port without a link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.
IfType	Specifies the type of OSPF interface (broadcast, NBMA, or passive). Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must use the same poll interval.
IfMtuIgnore	Specifies whether the VLAN ignores the MTU configuration. To allow the Virtual Services Platform 9000 to accept OSPF DD packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.
2. Double-click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **OSPF** tab.
The information on the OSPF tab applies only to a routed port or VLAN; that is, it uses an IP address.
7. To enable OSPF on the VLAN interface, select the **Enable** check box.
8. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.

9. To designate a router priority, in the **DesigRtrPriority** box, type the new value.
 10. To enable authentication, in the **AuthType** area, select either **simplePassword** or **MD5**.
 11. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
 12. Select the interface type you want to create.
 13. Click **Apply**.
-

Creating stubby or not-so-stubby OSPF areas

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA). Place stubby areas or NSSAs at the edge of an OSPF routing domain.

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

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

Table 14: Variable definitions

Variable	Value
Areald	Specifies a 32-bit integer that uniquely identifies an area. Area ID 0.0.0.0 is the OSPF backbone. For VLANs, keeping the default area setting on the interface causes LSDB inconsistencies.
ImportAsExtern	Specifies the method to import ASE link-state advertisements. The value can be importExternal (default), importNoExternal, or importNssa.
SpfRuns	Specifies the number of SPF calculations performed by OSPF.
AreaBdrRtrCount	Specifies the number of area border routers reachable within this area. Each SPF pass calculates this value, initially zero.

Variable	Value
AsBdrRtrCount	Specifies the number of autonomous system border routers reachable within this area. Each SPF pass calculates this value, initially zero.
AreaLsaCount	Specifies the total number of link state advertisements in this area LSDB, excluding AS-external LSAs.
AreaLsaCksumSum	Specifies the number of link-state advertisements. This sum excludes external (LS type 5) link-state advertisements. The sum determines if a change occurred in a router LSDB and compares the LSDB of two routers.
AreaSummary	Specifies whether to send summary advertisements in a stub area.
ActiveifCount	Specifies the number of active interfaces in this area.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Areas** tab.
The backbone ID has an area ID of 0.0.0.0.
 4. Click **Insert**.
 5. Configure the area ID.
 6. Select an option in the ImportAsExtern area.
To add a not-so-stubby (NSSA) area, select **importNssa**. To import external LSAs (create a normal OSPF area), select **importExternal**. To not import external LSAs (create a stubby area), select **importNoExternal**.
 7. Click **Apply**.
-

Configuring stub area metrics advertised by an ABR

Prerequisites

- Enable OSPF globally.
- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs.

For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 15: Variable definitions

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

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

Inserting OSPF area aggregate ranges

Prerequisites

- Enable OSPF globally.
- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 16: Variable definitions

Variable	Value
AreaID	Specifies the area in which the address exists.
LsdbType	Specifies the LSDB type: <ul style="list-style-type: none"> summaryLink—aggregated summary link nssaExternalLink—not so stubby area link
IP Address	Specifies the IP address of the network or subnetwork indicated by the range.
Mask	Specifies the network mask for the area range.
Effect	Specifies advertisement methods: <ul style="list-style-type: none"> advertiseMatching means advertise the aggregate summary LSA with the same LSID. doNotAdvertiseMatching means suppress all networks that fall within the entire range. advertiseDoNotAggregate means advertise individual networks.
AdvertiseMetric	Changes the advertised metric cost for the OSPF area range.

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

Enabling automatic virtual links

Prerequisites

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Configuring a manual virtual interface

Prerequisites

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 17: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Virtual If** tab.
 4. Click **Insert**.
 5. Specify the area ID of the transit area.
The transit area is the common area between two ABRs.
 6. Specify the neighbor ID.

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

7. Click **Insert**.
8. To verify that the virtual link is active, click **Refresh** and check the **State** column.
If the state is point-to-point, the virtual link is active. If the state is down, the virtual link configuration is incorrect.

Viewing virtual neighbors

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 18: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Virtual Neighbors** tab.
-

Configuring host routes

Prerequisites

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure host routes when the Avaya Virtual Services Platform 9000 resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

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

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

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

Table 19: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.

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

Enabling ASBR status

Prerequisites

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Enable the ASBR status to make the Avaya Virtual Services Platform 9000 an autonomous system boundary router (ASBR). Use ASBRs to advertise nonOSPF routes into OSPF domains so that the routes pass through the domain. A router can function as an ASBR if one or more of its interfaces connects to a non-OSPF network, for example, Routing Information Protocol (RIP), BGP, or Exterior Gateway Protocol (EGP).

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **General** tab.
 4. Select the **ASBdrRtrStatus** check box.
 5. Click **Apply**.
-

Managing OSPF neighbors

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

View or delete OSPF neighbors to control OSPF operations.

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

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

Viewing the link-state database

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 20: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Link State Database** tab.
-

Viewing the external link-state database

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

View the external LSDB to determine externally learned routing information.

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

Table 21: Variable definitions

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **Ext. Link State Database** tab.
-

Configuring route redistribution to OSPF

Prerequisites

- Enable OSPF globally.
 - Ensure that a route policy exists.
 - Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.
-

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



Important:

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

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

Table 22: Variable definitions

Variable	Value
DstVrfId	Specifies the destination virtual router forwarding instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrfId	Specifies the source VRF instance. You cannot configure this variable.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) to use for detailed redistribution of external routes from a specified source into an OSPF domain. Click the ellipsis (...) button and choose from the list in the dialog box.
Metric	Configures the OSPF route redistribution metric for basic redistribution. The value can be a range from 0–65535. A value of 0 indicates to use the original cost of the route.
MetricType	Configures the OSPF route redistribution metric type. The default is type 2. The cost of a type 2 route is the external cost, regardless of the interior cost. A type 1 cost is the sum of both the internal and external costs.
Subnets	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.

1. In the navigation tree, open the following folders: **Configuration, IP**.
2. Double-click **OSPF**.

3. Click the **Redistribute** tab.
 4. Click **Insert**.
 5. Select an option for the route source.
 6. Select the **enable** option button.
 7. Select a route policy.
 8. Configure the metric type.
 9. Configure the subnet.
 10. Click **Insert**.
-

Forcing shortest-path calculation updates

Prerequisites

Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

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

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **OSPF**.
 3. Click the **General** tab.
 4. In the **OspfAction** area, select the **runSpf** option button.
 5. Click **Apply**.
 6. Click **Yes** to force an SPF run.
After you initiate an SPF run, wait at least 10 seconds before you initiate another SPF run.
-

Chapter 7: OSPF configuration using ACLI

Configure Open Shortest Path First (OSPF) so that the Avaya Virtual Services Platform 9000 can use OSPF routing to communicate with other OSPF routers and to participate in OSPF routing.

- [Configuring OSPF globally](#) on page 83
- [Configuring OSPF for a port or VLAN](#) on page 86
- [Viewing OSPF errors on a port](#) on page 88
- [Configuring OSPF areas on the router](#) on page 89
- [Configuring OSPF aggregate area ranges on the router](#) on page 90
- [Enabling automatic virtual links](#) on page 92
- [Configuring an OSPF area virtual interface](#) on page 92
- [Configuring an OSPF area on a VLAN or port](#) on page 94
- [Configuring an OSPF host route](#) on page 96
- [Configuring OSPF NBMA neighbors](#) on page 97
- [Applying OSPF route acceptance policies](#) on page 98
- [Viewing the OSPF link-state database](#) on page 100
- [Viewing the OSPF external link-state database](#) on page 100
- [Configuring route redistribution to OSPF](#) on page 101
- [Configuring interVRF route redistribution for OSPF](#) on page 103
- [Forcing shortest-path calculation updates](#) on page 104

Configuring OSPF globally

Configure OSPF parameters on the switch to control how OSPF behaves on the system. The Avaya Virtual Services Platform 9000 uses global parameters to communicate with other OSPF routers. Globally configure OSPF before you configure OSPF for an interface, port, or VLAN.

Prerequisites

- Ensure that the Virtual Services Platform 9000 has an IP address.
- You must log on to the OSPF Router Configuration mode in ACLI to configure the GlobalRouter.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf` to

commands. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure steps

-
1. Configure the OSPF router ID:
`router-id {A.B.C.D}`
 2. Configure the router as an autonomous system boundary router (ASBR):
`as-boundary-router enable`
 3. Enable the automatic creation of OSPF virtual links:
`auto-vlink`
 4. Configure the OSPF default metrics:
`default-cost [{ethernet|fast-ethernet|gig-ethernet|ten-gig-ethernet} <1-65535>]`
 5. Configure the OSPF hold-down timer value:
`timers basic holddown <3-60>`
 6. Enable the RFC1583 compatibility mode:
`rfc1583-compatibility enable`
 7. Enable the router to issue OSPF traps:
`trap enable`
 8. Verify the OSPF configuration:
`show ip ospf [vrf WORD<0-16>] [vrfids WORD<0-512>]`
 9. Exit OSPF Router Configuration mode:
`exit`
You return to Global Configuration mode.
 10. Enable OSPF for the switch:
`router ospf enable`
-

Result

Variable definitions

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

Table 23: Variable definitions

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

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

Table 24: Variable definitions

Variable	Value
<code>ethernet <1-65535></code>	Configures the OSPF default metrics. The range is 1–65535. ethernet is for 10 Mb/s Ethernet (default is 100).
<code>fast-ethernet <1-65535></code>	Configures the OSPF default metrics. The range is 1–65535. fast-ethernet is for 100 Mb/s (Fast) Ethernet (default is 10).
<code>gig-ethernet <1-65535></code>	Configures the OSPF default metrics. The range is 1–65535. gig-ethernet is for Gigabit Ethernet (default is 1).
<code>ten-gig-ethernet <1-65535></code>	Configures the OSPF default metrics. The range is 1–65535. ten-gig-ethernet is for 10 Gigabit Ethernet (default is 1).

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

Table 25: Variable definitions

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

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

Table 26: Variable definitions

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

Configuring OSPF for a port or VLAN

Prerequisites

- Enable OSPF globally.
- Ensure that the VLAN exists.
- Ensure that the port or VLAN uses an IP address.
- Ensure that you know the network OSPF password to use password authentication or that you know the Message Digest 5 (MD5) key to use MD5 authentication.
- You must log on to the Interface Configuration mode for the port or VLAN in ACLI.

Configure OSPF parameters on a port or VLAN to control how OSPF behaves on the port or VLAN.

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

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

Table 27: Variable definitions

Variable	Value
<code>advertise-when-down</code> <code>enable</code>	Enables or disables AdvertiseWhenDown. If enabled, OSPF advertises the network on this interface as up, even if the port is down. The default is disabled. After you configure a port with no link and enable <code>advertise-when-down</code> , OSPF does not advertise the route until the port is active. OSPF advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
<code>area {A.B.C.D}</code>	Configures the OSPF identification number for the area, typically formatted as an IP address.
<code>authentication-key</code> <code>WORD<0-8></code>	Configures the eight-character simple password authentication key for the port or VLAN.
<code>authentication-type</code> <code><message-digest </code> <code>none simple></code>	Configures the OSPF authentication type for the port: none, simple password, or MD5 authentication. If simple, all OSPF updates the interface receives must contain the authentication key specified by the <code>area authentication-key</code> command. If MD5, they must contain the MD5 key.
<code>cost <0-65535></code>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.

Variable	Value
dead-interval <0-2147483647>	Configures the router OSPF dead interval—the number of seconds the OSPF neighbors of a switch must wait before they assume the OSPF router is down. The default is 40. The value must be at least four times the hello interval.
enable	Enables OSPF on the port or VLAN.
hello-interval <1-65535>	Configures the OSPF hello interval, which is the number of seconds between hello packets sent on this interface. The default is 10.
message-digest-key <1-255> md5 WORD<0-16>	Configures the MD5 key. At most, you can configure two MD5 keys for an interface. <1-255> is the ID for the message digest key WORD<0-16> is an alphanumeric password of up to 16 bytes {string length 0-16}
mtu-ignore enable	Enables maximum transmission unit (MTU) ignore. To allow the Virtual Services Platform 9000 to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
network <broadcast nbma passive>	Specifies the type of OSPF interface.
poll-interval <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
primary-md5-key <1-255>	Changes the primary key used to encrypt outgoing packets. <1-255> is the ID for the new message digest key.
priority <0-255>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you configure the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
retransmit-interval <0-3600>	Configures the retransmit interval for the virtual interface, the number of seconds between link-state advertisement retransmissions.
transit-delay <0-3600>	Configures the transit delay for the virtual interface, which is the estimated number of seconds required to transmit a link-state update over the interface.
vlan <1-4084>	Specifies the VLAN ID. This variable applies only to VLAN interfaces, not ports.

1. Configure the OSPF interface area ID:

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

2. Enable OSPF routing:

```
ip ospf enable
```

3. Choose the OSPF update authentication method:

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

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

4. If you choose simple, you must configure the password. If you choose MD5, you must configure the MD5 key:

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

OR

```
ip ospf message-digest-key <1-255> md5 WORD<0-16>
```

5. Specify the interface type:

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

6. Configure the remaining parameters as required, or accept their default values.

Viewing OSPF errors on a port

Prerequisites

You must log on to Privileged EXEC mode in ACLI.

Check OSPF errors for administrative and troubleshooting purposes.

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

Table 28: Variable definitions

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
vrf WORD<0-16>	Specifies the VRF by name.
vrfids WORD<0-512>	Specifies a range of VRFs by ID number.

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


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

Configuring OSPF areas on the router

Prerequisites

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

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA). Place stubby or NSSAs at the edge of an OSPF routing domain.

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

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

Table 29: Variable definitions

Variable	Value
default-cost <0-16777215>	Specifies the stub area default metric for this stub area, which is the cost from 0–16777215. This metric value applies at the indicated type of service.
import <external noexternal nssa>	Specifies the type of area: <ul style="list-style-type: none"> • external—stub and NSSA are both false • noexternal—configures the area as stub area. • nssa—configures the area as NSSA.
import-summaries enable	Configures the area support to import summary advertisements into a stub area. Use this variable only if the area is a stub area.
stub	Configures the import external option for this area as stub. A stub area has only one exit point (router interface) from the area.

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

Table 30: Variable definitions

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

-
1. Create an OSPF area:

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

2. Specify the area type:

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

3. Configure other OSPF area parameters as required.

4. Ensure that the configuration is correct:

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

Configuring OSPF aggregate area ranges on the router

Prerequisites

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

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

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

Table 31: Variable definitions

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

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

Table 32: Variable definitions

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

1. Configure an OSPF area range:

```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
```
2. Configure the advertised metric cost:

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

```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
advertise-mode <summarize|suppress|no-summarize>
```
4. Ensure that the configuration is correct:

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

Enabling automatic virtual links

Prerequisites

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

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

Enable the automatic virtual links feature for the router:

```
auto-vlink
```

Configuring an OSPF area virtual interface

Prerequisites

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

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

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

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

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

Table 33: Variable definitions

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the area ID and the virtual interface ID.
authentication-key WORD<0-8>	Configures the authentication key of up to eight characters.
authentication-type <message-digest none simple>	Configures the authentication type for the OSPF area. <i>auth-type</i> is none, simple password, or MD5 authentication. If simple, all OSPF updates received by the interface must contain the authentication key specified by the area authentication-key command. If MD5, they must contain the MD5 key. The default is none.
dead-interval <0-2147483647>	Configures the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
hello-interval <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
primary-md5-key <1-255>	Changes the primary key that encrypts outgoing packets. <1-255> is the ID for the message digest key.
retransmit-interval <0-3600>	Configures the retransmit interval for the virtual interface, the number of seconds between LSA retransmissions. The range is from 1–3600.
transit-delay <0-3600>	Configures the transit delay for the virtual interface, the estimated number of seconds required to transmit a link-state update over the interface. The range is from 1–3600.

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

Table 34: Variable definitions

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the area ID and the virtual interface ID.
<1-255>	Specifies the ID for the message digest key
md5-key WORD<1-16>	Adds an MD5 key to the interface. At most, you can configure two MD5 keys to an interface. WORD<1-16> is an alphanumeric password of up to 16 characters.

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

Table 35: Variable definitions

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

1. Create an OSPF area virtual interface:

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

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

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

3. Configure optional parameters, as required.

4. Ensure that the configuration is correct:

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

Configuring an OSPF area on a VLAN or port

Prerequisites

- Enable OSPF globally.
- Ensure that the VLAN exists.
- You must log on to Interface Configuration mode in ACLI for the VLAN or port.

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

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

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

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

Table 36: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the area ID.
authentication-key WORD<0-8>	Configures the eight-character simple password authentication key for the port or VLAN.
authentication-type <message-digest none simple>	Configures the authentication type for the OSPF area. The type is none, simple password, or MD5 authentication. If simple, all OSPF updates received by the interface must contain the authentication key specified by the area authentication-key command. If MD5, they must contain the MD5 key. The default is none.
cost <0-65535>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.
dead-interval <0-2147483647>	Configures the the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
hello-interval <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
mtu-ignore enable	Enables MTU ignore. To allow the Virtual Services Platform 9000 to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
network <broadcast nbma passive>	Specifies the type of OSPF interface.
poll-interval <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
primary-md5-key <1-255>	Changes the primary key used to encrypt outgoing packets. <1-255> is the ID for the message digest key.
priority <0-255>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you set the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
retransmit-interval <0-3600>	Configures the retransmit interval: the number of seconds between LSA retransmissions. The range is from 1–3600.

Variable	Value
transit-delay <0-3600>	Configures the transit delay: the estimated number of seconds it takes to transmit a link-state update over the interface. The range is from 1–3600.

-
1. Create an OSPF area on the VLAN or port:
`ip ospf area {A.B.C.D}`
 2. Specify the type of network:
`ip ospf network <broadcast|nbma|passive>`
 3. Configure other OSPF area parameters as required.
-

Configuring an OSPF host route

Prerequisites

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

Configure host routes when the Avaya Virtual Services Platform 9000 resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

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

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

Use the data in the following table to use the `host-route` commands.

Table 37: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the IP address of the host router in a.b.c.d format.

Variable	Value
<code>metric <0-65535></code>	Configures the metric (cost) for the host route.

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

Table 38: Variable definitions

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

-
1. Create a host route:

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

2. Ensure that the configuration is correct:

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

Configuring OSPF NBMA neighbors

Prerequisites

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

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

Use the data in the following table to see the `neighbor` command.

Table 39: Variable definitions

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

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

Table 40: Variable definitions

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

-
1. Create an NBMA OSPF neighbor:

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

2. Ensure that the configuration is correct:

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

Applying OSPF route acceptance policies

Prerequisites

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

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

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

Table 41: Variable definitions

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

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

Table 42: Variable definitions

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

-
1. Create an acceptance policy instance:

```
accept adv-rtr {A.B.C.D}
```
 2. Configure the type of metric to accept:

```
accept adv-rtr {A.B.C.D} metric-type <type1|type2|any>
```
 3. Indicate the route policy:

```
accept adv-rtr {A.B.C.D} route-policy WORD<0-64>
```
 4. Enable a configured OSPF route acceptance instance:

```
accept adv-rtr {A.B.C.D} enable
```
 5. Ensure that the configuration is correct:

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

Viewing the OSPF link-state database

Prerequisites

You must log on to Privileged EXEC mode in ACLI.

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

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

Table 43: Variable definitions

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

View the OSPF link-state database:

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

Viewing the OSPF external link-state database

Prerequisites

You must log on to Privileged EXEC mode in ACLI.

View the LSDB to determine externally learned routing information.

Information appears for all metric types or for the type you specify.

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

Table 44: Variable definitions

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

View the OSPF autonomous system external (ASE) link-state advertisements:

```
show ip ospf ase [metric-type <1-2>] [vrf WORD<0-16>] [vrfrids
WORD<0-512>]
```

Configuring route redistribution to OSPF

Prerequisites

- Enable OSPF globally.
- Ensure that a route policy exists.
- You must log on to OSPF Router Configuration mode in ACLI.

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

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

Table 45: Variable definitions

Variable	Value
<code>enable</code>	Enables the OSPF route redistribution instance.
<code>metric <0-65535></code>	Configures the metric to apply to redistributed routes.
<code>metric-type <type1 type2></code>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.

Variable	Value
route-policy WORD<0-64>	Configures the route policy to apply to redistributed routes.
subnets <allow suppress>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
vrf-src WORD<0-16>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
WORD<0-32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, direct, rip, ospf, or static.

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

Table 46: Variable definitions

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

1. Create the redistribution instance:

```
redistribute WORD<0-32> <ospf|bgp|static|direct|rip> [vrf-  
src WORD<0-16>]
```

2. Apply a route policy if required:

```
redistribute WORD<0-32> route-policy WORD<0-64> [vrf-src  
WORD<0-16>]
```

3. Configure other parameters, as required.

4. Enable the redistribution.

```
redistribute WORD<0-32> enable [vrf-src WORD<0-16>]
```

5. Ensure that the configuration is correct:

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

6. Exit OSPF Router Configuration mode.

```
exit
```

You are now in Global Configuration mode.

7. Apply the redistribution.

```
ip ospf apply redistribute WORD<0-32> [vrf WORD<0-16>] [vrf-
src WORD<0-16>]
```

Changes do not take effect until you apply them.

Configuring interVRF route redistribution for OSPF

Prerequisites

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.
- You must log on to VRF Router Configuration mode in ACLI.

Use route redistribution so that a VRF interface can announce routes learned by other protocols, for example, OSPF or BGP. The Avaya Virtual Services Platform 9000 supports interVRF route redistribution. Use a route policy to control the redistribution of routes.

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

Table 47: Variable definitions

Variable	Value
enable	Enables the OSPF route redistribution instance.
metric <0-65535>	Configures the metric to apply to redistributed routes.
metric-type <type1 type2 any>	Specifies a metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
route-policy WORD<0-64>	Configures the route policy to apply to redistributed routes.
subnets <allow suppress>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
vrf-src WORD<0-16>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
WORD<0-32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, direct, rip, ospf, or static.

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

Table 48: Variable definitions

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

1. Create the redistribution instance:

```
ip ospf redistribute WORD <1-32>
```

2. Apply a route policy if required:

```
ip ospf redistribute WORD <1-32> route-policy WORD<0-64>
[vrf-src WORD<0-16>]
```

3. Configure other parameters, as required.

4. Enable the redistribution:

```
ip ospf redistribute WORD <1-32> enable [vrf-src WORD<0-16>]
```

5. Ensure that the configuration is correct:

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

6. Exit VRF Router Configuration mode.

```
exit
```

You are now in Global Configuration mode.

7. Apply the redistribution:

```
ip ospf apply redistribute WORD<0-32> [vrf WORD<0-16>] [vrf-
src WORD<0-16>]
```

Forcing shortest-path calculation updates

Prerequisites

You must log on to Privileged EXEC mode in ACLI.

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

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

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

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

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by name.

Force the router to update its shortest-path calculations:

```
ip ospf spf-run [vrf WORD<0-16>]
```

Chapter 8: RIP configuration using EDM

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use Enterprise Device Manager (EDM) to configure and manage the RIP on an Avaya Virtual Services Platform 9000.

- [Configuring RIP globally](#) on page 107
- [Configuring RIP interface compatibility](#) on page 108
- [Configuring RIP on an interface](#) on page 111
- [Configuring RIP on a port](#) on page 113
- [Configuring RIP on a VLAN](#) on page 115
- [Configuring route redistribution to RIP](#) on page 117

Configuring RIP globally

Prerequisites

Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure RIP global parameters on the switch to control how RIP behaves on the system.

In the Avaya Virtual Services Platform 9000, all router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

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

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

Table 49: Variable definitions

Variable	Value
Operation	Enables or disables RIP on all interfaces. The default is disabled.

Variable	Value
UpdateTime	Specifies the time interval between RIP updates for all interfaces. The default is 30 seconds, and the range is 0–360.
RouteChanges	Specifies the number of route changes RIP made to the IP route database. RouteChanges does not include the refresh of a route age.
Queries	Specifies the number of responses sent to RIP queries received from other systems.
HoldDownTime	Configures the length of time that RIP continues to advertise a network after the network is unreachable. The range is 0–360 seconds. The default is 120 seconds.
TimeOutInterval	Configures the RIP timeout interval. The range is 15–259200 seconds. The default is 180 seconds.
DefImportMetric	Configures the default import metric used to import a route into a RIP domain. To announce OSPF internal routes into a RIP domain, if the policy does not specify a metric, you must use the default import metric. OSPF external routes use the external cost. The range is 0–15 and the default is 8.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **RIP**.
 3. Click the **Globals** tab.
 4. Select the **enable** option button.
 5. Configure other global RIP parameters as required.
 6. Click **Apply**.
-

Configuring RIP interface compatibility

Prerequisites

- Configure a routing interface (either a router port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIP globally.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non0 VRFs.

For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure RIP parameters on an interface to control how RIP behaves on the interface. You can specify the RIP version to use on interfaces that you configure to send (supply) or receive (listen to) RIP updates.

On an interface, RIP does not operate until you enable it globally and on the interface.

Although visible, Avaya Virtual Services Platform 9000 does not support the AuthType and AuthKey parameters.

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

Table 50: Variable definitions

Variable	Value
Address	Specifies the IP address of the router interface.
Domain	Specifies the value inserted into the routing domain parameter of all RIP packets sent on this interface.
AuthType	Specifies the type of authentication to use on this interface.
AuthKey	Specifies the authentication key whenever AuthType is not noAuthentication.
Send	<p>Specifies the update version the router sends on this interface:</p> <ul style="list-style-type: none"> • DoNotSend—no RIP updates sent on this interface • ripVersion1—RIP updates compliant with RFC1058 • rip1Compatible—broadcast RIPv2 updates using RFC1058 route subassumption rules • ripVersion2—multicast RIPv2 updates <p>The default is rip1compatible.</p>
Receive	<p>Indicates which versions of RIP updates to accept:</p> <ul style="list-style-type: none"> • rip1 • rip2 • rip1OrRip2 <p>The default is rip1OrRip2. Rip2 and rip1OrRip2 imply receipt of multicast packets.</p>

1. In the navigation tree, open the following folders: **Configuration, IP**.
2. Double-click **RIP**.
3. Click the **Interface** tab.

4. Double-click the **Send** value to edit it, and then select the RIP version datagrams the router sends.
 5. Double-click the **Receive** value to edit it, and then select the RIP version datagrams for which the router listens.
 6. Click **Apply**.
-

Job aid

Choose one of three options for receiving RIP updates:

- rip1OrRip2—accepts RIPv1 or RIPv2 updates
- rip1—accepts RIPv1 updates only
- rip2—accepts RIPv2 updates only

The following table describes the four RIP send modes that Virtual Services Platform 9000 supports. You can configure RIP send modes on all router interfaces.

Table 51: RIP send modes

Send mode	Description	Result
rip1comp	Broadcasts RIPv2 updates using RFC1058 route consumption rules. This mode is the default mode on the Virtual Services Platform 9000.	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.1.2.255) • RIP update is formed as a RIP-2 update, including network mask • RIP version = 2
rip1	Broadcasts RIP updates that are compliant with RFC1058	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.1.2.255) • RIP update is formed as a RIP-1 update, no network mask included • RIP version = 1
rip2	Broadcasts multicast RIPv2 updates	<ul style="list-style-type: none"> • Destination MAC is a multicast, 01-00-5e-00-00-09 • Destination IP is the RIP-2 multicast address, 224.0.0.9

Send mode	Description	Result
		<ul style="list-style-type: none"> • RIP update is formed as a RIP-2 update including network mask • RIP version = 2
nosend	Does not send RIP updates on the interface	None

Configuring RIP on an interface

Prerequisites

Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure RIP parameters to control and optimize RIP routing for the interface.

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

Table 52: Variable definitions

Variable	Value
Address	Shows the address of the entry in the IP RIP interface table.
Interface	Indicates the index of the RIP interface.
Enable	Shows if the RIP interface is enabled or disabled.
Supply	Enables (true) or disables (false) the ability to send RIP updates on this interface.
Listen	Configures whether the switch learns routes on this interface.
Poison	Configures whether to advertise RIP routes learned from a neighbor back to the neighbor. If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, RIP poisons the RIP updates, sent to the neighbor from which a route is learned, with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network.
DefaultSupply	Enables (true) or disables (false) an advertisement of a default route on this interface. This command takes effect only if a default route exists in the routing table.

Variable	Value
DefaultListen	Enables (true) or disables (false) the switch to accept the default route learned through RIP on this interface. The default is disabled. Enable DefaultListen to add a default route to the route table if another route advertises it.
TriggeredUpdate	Enables (true) or disables (false) the switch to send RIP updates from this interface.
AutoAggregate	Enables (true) or disables (false) automatic route aggregation on this interface. If enabled, the switch automatically aggregates routes to their natural mask when an interface advertises them. The default is disabled.
InPolicy	Determines if RIP can learn routes on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if RIP advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The range is 1–15. The default is 1.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **RIP**.
 3. Click the **Interface Advance** tab.
 4. Double-click a RIP parameter to edit it, as required.
 5. Click **Apply**.
-

Job aid

The following table indicates the relationship between switch action and the RIP supply and listen settings.

Table 53: RIP supply and listen settings and switch action

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
Disabled	Disabled			Sends no RIP updates.

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
Enabled	Disabled			Sends RIP updates except the default.
Disabled	Enabled			Sends only the default (default route must exist in routing table).
Enabled	Enabled			Sends RIP updates including the default route (if it exists).
		Disabled	Disabled	Does not listen to RIP updates.
		Enabled	Disabled	Listens to all RIP updates except the default.
		Disabled	Enabled	Listens only to the default.
		Enabled	Enabled	Listens to RIP updates including the default route (if it exists).

Configuring RIP on a port

Prerequisites

- Assign an IP address to the port.
- Configure RIP and enable it globally.

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

- Enable RIP on the interface.

Configure RIP on a port so that the port can participate in RIP routing.

On an interface, RIP does not operate until you enable it globally and on the interface.

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

Table 54: Variable definitions

Variable	Value
Enable	Enables or disables RIP on the port.

Variable	Value
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the port is down. The default is false. If you configure a port with no link and enable AdvertiseWhenDown, the port does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if this interface advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.

Variable	Value
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

-
1. In the Device Physical View tab, select a port.
 2. In the navigation tree, open the following folders: **Configuration, Edit, Port**.
 3. Double-click **IP**.
 4. Click the **RIP** tab.
 5. Configure the RIP parameters as required.
 6. Click **Apply**.
-

Configuring RIP on a VLAN

Prerequisites

- Configure the VLAN.
- Assign an IP address to the VLAN.
- Enable RIP globally.
- Enable RIP on the interface.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

Configure RIP on a VLAN so that the VLAN acts as a routed VLAN (a virtual router).

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

Table 55: Variable definitions

Variable	Value
Enable	Enables or disables RIP on the VLAN.
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the interface is down. The default is false. If you configure a VLAN with no link and enable AdvertiseWhenDown, the VLAN does not advertise the route until the VLAN is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.

Variable	Value
OutPolicy	Determines if this interface advertises a route from the routing table. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

-
1. In the navigation tree, open the following folders: **Configuration, VLAN**.
 2. Double-click **VLANS**.
 3. Click the **Basic** tab.
 4. Select a VLAN.
 5. Click **IP**.
 6. Click the **RIP** tab.
 7. Configure the VLAN RIP parameters as required.
 8. Click **Apply**.
-

Configuring route redistribution to RIP

Prerequisites

- Enable RIP globally.
- Configure a route policy.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non0 VRFs. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

**Important:**

Changing the RIP redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Avaya recommends that if you want to change default preferences for a RIP redistribute context, you must do so before you enable the protocols.

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

Table 56: Variable definitions

Variable	Value
DstVrfId	Specifies the destination VRF instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrfId	Specifies the source VRF instance. You cannot configure this variable.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) a RIP redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) that redistributes external routes from a specified source into an RIP domain. Click the ellipsis (...) button and choose from the list in the Route Policy dialog box.
Metric	Configures the RIP route redistribution metric for basic redistribution. The value can be in the range 0–65535. A value of 0 indicates to use the original cost of the route.

-
1. In the navigation tree, open the following folders: **Configuration, IP**.
 2. Double-click **RIP**.
 3. Click the **Redistribute** tab.
 4. Click **Insert**.
 5. Configure the source of the routes to redistribute.
 6. Select **enable**.
 7. Select the route policy to apply to redistributed routes.
 8. Configure a metric value.
 9. Click **Insert**.
-

Chapter 9: RIP configuration using ACLI

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use the Avaya command line interface (ACLI) to configure and manage RIP on an Avaya Virtual Services Platform 9000.

- [Configuring RIP globally](#) on page 119
- [Configuring RIP on an interface](#) on page 121
- [Configuring route redistribution to RIP](#) on page 124
- [Configuring interVRF route redistribution for RIP](#) on page 125
- [Forcing a RIP update for a port or VLAN](#) on page 127

Configuring RIP globally

Prerequisites

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

Configure RIP parameters on the switch to control how RIP behaves on the system.

In the Avaya Virtual Services Platform 9000, all router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

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

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

Table 57: Variable definitions

Variable	Value
<code>default-metric</code> <0-15>	Configures the value of default import metric to import a route into a RIP domain. To announce OSPF internal routes into RIP domain, if the policy does not specify a metric value, the default is used. For OSPF external routes, the external cost is used. The default is 8.

Variable	Value
<code>domain <0-39321></code>	Specifies the RIP domain from 0–39321. The default is 0.
<code>holddown <0-360></code>	Configures the RIP hold-down timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.
<code>network {A.B.C.D}</code>	Enables RIP on an IP network.
<code>timeout <15-259200></code>	Configures the RIP timeout interval. The default is 180.
<code>update <1-360></code>	Configures the RIP update timer. The update time is the time interval, in seconds, between RIP updates. The default is 30.

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

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

-
1. Define the default-import-metric for the switch:
`default-metric <0-15>`
 2. Optionally, configure one or more timer values:
`timers basic timeout <15-259200> [holddown <0-360>] [update <1-360>]`
 3. Specify the RIP domain:
`domain <0-39321>`
 4. Enable RIP on an IP network:
`network {A.B.C.D}`
 5. Exit RIP Router Configuration mode.
`exit`
You are now in Global Configuration mode.
 6. After the configuration is complete, enable RIP globally:
`router rip enable`
 7. Check that your configuration is correct:
`show ip rip [vrf WORD<0-16>] [vrfids WORD<0-512>]`
-

Configuring RIP on an interface

Prerequisites

- Assign an IP address to the port or VLAN.
- Configure RIP and enable it globally.
- Configure in and out policies.
- You must log on to Interface Configuration mode in ACLI.

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

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

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

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

Table 58: Variable definitions

Variable	Value
<code>advertise-when-down</code> <code>enable</code>	Enables or disables AdvertiseWhenDown. If enabled, RIP advertises the network on this interface as up, even if the port is down. The default is disabled. If you configure a port with no link and enable advertise-when-down, it does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
<code>auto-aggregation</code> <code>enable</code>	Enables or disables automatic route aggregation on the port. If enabled, the switch automatically aggregates routes to their natural mask when an interface in a different class network advertises them. The default is disable.
<code>cost <1-15></code>	Configures the RIP cost for this port (link).
<code>default-listen</code> <code>enable</code>	Enables DefaultListen. The switch accepts the default route learned through RIP on this interface. The default is disabled.
<code>default-supply</code> <code>enable</code>	Enables DefaultSupply. If enabled, this interface must advertise a default route. The default is false. RIP advertises the default route only if it exists in the routing table.
<code>enable</code>	Enables RIP routing on the port.

Variable	Value
<code>holddown <0-360></code>	Configures the RIP holddown timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.
<code>in-policy WORD<0-64></code>	Configures the policy name for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when RIP adds it to the routing table.
<code>listen enable</code>	Specifies that the routing switch learns RIP routes through this interface. If enabled, the switch listens for a default route without listening for all routes. The default is enable.
<code>out-policy WORD<0-64></code>	Configures the policy name for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement. <i>policy name</i> is a string of length 0–64 characters.
<code>poison enable</code>	Enables Poison Reverse. If you disable Poison Reverse (no poison enable), Split Horizon is enabled. By default, Split Horizon is enabled. If you enable Split Horizon, the interface does not advertise IP routes learned from an immediate neighbor back to the neighbor. If you enable Poison Reverse, the RIP updates sent to a neighbor from which a route is learned are poisoned with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network. These mechanisms prevent routing loops.
<code>port {slot/port[- slot/port][,...]}</code>	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
<code>receive version <rip1 rip2 rip1orrip2></code>	Indicates which RIP update version to accept on this interface. The default is rip1orrip2.
<code>send version <notsend rip1 rip1comp rip2></code>	Indicates which RIP update version the router sends from this interface. ripVersion1 implies sending RIP updates that comply with RFC1058. rip1comp implies broadcasting RIP2 updates using RFC1058 route subassumption rules. The default is rip1Compatible.
<code>supply enable</code>	Specifies that the switch advertises RIP routes through the port. The default is enable.
<code>timeout <15-259200></code>	Configures the RIP timeout interval in seconds. The default is 180.
<code>triggered enable</code>	Enables automatic triggered updates for RIP.

1. Define the cost:
`ip rip cost <1-15>`
2. Specify an in policy for filtering inbound RIP packets:
`ip rip in-policy WORD<0-64>`
3. Specify an out policy for filtering outbound RIP packets:
`ip rip out-policy WORD<0-64>`
4. Enable RIP:
`ip rip enable`
5. Specify the send mode:
`ip rip send version <notsend|rip1|rip1comp|rip2>`
6. Specify the receive mode:
`ip rip receive version <rip1|rip2|rip1orrip2>`
7. Change other RIP parameters from their default values as required.

Example of configuring RIP on an interface

This configuration example shows how to configure Virtual Services Platform (R1 in the following figure) to operate only in RIP version 2 mode.




Figure 15: Configuration example-RIPv2 only

1. Enable RIPv2 send mode on VLAN 2:
`VSP-9012:1(config-if) # ip rip send version rip2`
2. Enable RIPv2 receive mode on VLAN 2:

```
VSP-9012:1(config-if)# ip rip receive version rip2
```

3. Repeat these commands on VLAN 3 and the port interfaces.

Configuring route redistribution to RIP

Prerequisites

- Enable RIP globally.
- Configure a route policy.
- You must log on to RIP Router Configuration mode in ACLI.

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

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

Table 59: Variable definitions

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

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

Table 60: Variable definitions

Variable	Value
<code>vrf WORD<0-16></code>	Specifies the VRF instance.
<code>vrfids WORD<1-512></code>	Specifies a range of VRF IDs.

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

Table 61: Variable definitions

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

1. Create the redistribution instance:

```
redistribute WORD<0-32> [vrf-src WORD<0-16>]
```

2. Apply a route policy, if required:

```
redistribute WORD<0-32> route-map WORD<0-64> [vrf-src WORD<0-16>]
```

3. Configure other parameters.

4. Enable the redistribution:

```
redistribute WORD<0-32> enable [vrf-src WORD<0-16>]
```

5. Ensure that the configuration is correct:

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

6. Exit RIP Router Configuration mode.

```
exit
```

You are now in Global Configuration mode.

7. Apply the redistribution:

```
ip rip apply redistribute WORD<0-32> [vrf WORD<0-16>] [vrf-src WORD<0-16>]
```

Configuring interVRF route redistribution for RIP

Prerequisites

- Enable RIP globally.
- Configure a route policy.

- Configure the VRFs.
- You must log on to VRF Router Configuration mode in ACLI.

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

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

Table 62: Variable definitions

Variable	Value
WORD<0-32>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, direct, rip, ospf, or static.
vrf-src WORD<0-16>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
metric <0-65535>	Configures the metric to apply to redistributed routes.
route-map <WORD 0-64>	Configures the route policy to apply to redistributed routes.

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

Table 63: Variable definitions

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

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

Table 64: Variable definitions

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

1. Create the redistribution instance:

```
ip rip redistribute WORD<0-32>
```

2. Apply a route policy, if required:

```
ip rip redistribute WORD<0-32> route-map WORD<0-64> [vrf-src  
WORD<0-16>]
```

3. Configure other parameters.

4. Enable the redistribution:

```
ip rip redistribute WORD<0-32> enable [vrf-src WORD<0-16>]
```

5. Ensure that the configuration is correct:

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

6. Exit VRF Router Configuration mode:

```
exit
```

You are now in Global Configuration mode.

7. Apply the redistribution:

```
ip rip apply redistribute WORD<0-32> [vrf WORD<0-16>] [vrf-  
src WORD<0-16>]
```

Forcing a RIP update for a port or VLAN

Prerequisites

You must log on to Interface Configuration mode in ACLI.

Force RIP to update the routing table so that the port or VLAN uses the latest routing information.

If you perform this procedure, you also update the tables for all VRFs associated with the port or VLAN.

1. Update the routing table for a port:

```
action triggerRipUpdate
```

2. Update the routing table for a VLAN:

```
ip rip triggered enable
```

Chapter 10: Common procedures using EDM

The following section describes common procedures that you use while you configure Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) on the Avaya Virtual Services Platform 9000 using Enterprise Device Manager (EDM).

[Configuring interVRF route redistribution policies](#) on page 129

Configuring interVRF route redistribution policies

Prerequisites

- VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required. For information about how to log on to EDM for a non0 VRF, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals, NN46250-103*.

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

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

Table 65: Variable definitions

Variable	Value
DstVrflid	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrflid	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine

Variable	Value
	whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
Subnets	Indicates that subnets must be advertised individually (applies to OSPF only).

-
1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Double-click **Policy**.
 3. Click the **Route Redistribution** tab.
 4. Click **Insert**.
 5. Click the ellipsis (...) button near the **DstVrfId** box to select the source and destination VRF IDs.
 6. Click the ellipsis (...) button near the **SrcVrfId** box to select the source and destination VRF IDs.
 7. In the **Protocol** option box, select the protocol.
 8. In the **RouteSource** option box, select the route source.
 9. Select **enable**.
 10. Click the ellipsis (...) button near the **RoutePolicy** box to choose the route policy to apply to the redistributed routes.
 11. Configure other parameters as required.
 12. Click **Insert**.
 13. Click the **Applying Policy** tab.
 14. Select **RedistributeApply**.
 15. Click **Apply**.
-

Chapter 11: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Navigation

- [Getting technical documentation](#) on page 131
- [Getting product training](#) on page 131
- [Getting help from a distributor or reseller](#) on page 131
- [Getting technical support from the Avaya Web site](#) on page 132

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Index

Numerics

10000MbpsPortDefaultMetric	54
1000MbpsPortDefaultMetric	54
100MbpsPortDefaultMetric	54
10MbpsPortDefaultMetric	54

A

accept policy	19
ActiveifCount	69
Address	108 , 111
AddressLessIf	57 , 62
AddressLessIndex	61
AdminStat	54 , 57
Advertisement	79
AdvertiseMetric	71
AdvertiseWhenDown	64 , 66 , 113 , 115
Age	78 , 79
alternative routes	23
announce policy	19
Area	75
AreaBdrRtrCount	69
AreaBdrRtrStatus	54
Areald	57 , 64 , 66 , 69 , 70 , 78
AreaID	71 , 73 , 76
AreaLsaChecksumSum	69
AreaLsaCount	69
AreaSummary	69
AsBdrRtrCount	69
ASBdrRtrStatus	54
ASBR	
enabling using EDM	77
authentication key	92
authentication type	92 , 94
AuthKey	57 , 64 , 66 , 73 , 108
AuthType	57 , 64 , 66 , 73 , 108
AutoAggregate	111
AutoAggregateEnable	113 , 115
automatic route aggregation	111 , 121
AutoVirtLinkEnable	54

B

black hole static route	17
brouter port	16

bridging traffic	16
description	16
IP routing	16

C

Checksum	78 , 79
CIDR	12
Circuitless IP, see CLIP	18
CLIP	18
Cost	111 , 113 , 115
customer service	131

D

dead interval	92 , 94
OSPF	94
DefaultListen	111 , 113 , 115
DefaultSupply	111 , 113 , 115
DefImportMetric	107
designated router	32
backup	32
DesigRtrPriority	64 , 66
Dijkstras algorithm	28
distributor	131
documentation	131
Domain	108
DstVrflid	80 , 117 , 129

E

ECMP	
description	23
Effect	71
Enable	64 , 66 , 80 , 111 , 113 , 115 , 117 , 129
Equal Cost MultiPath. See ECMP	23
errors	88
Events	61 , 73 , 75
ExternLsaChecksumSum	54
ExternLsaCount	54

H

Hello protocol	28 , 32
HelloInterval	57 , 64 , 66 , 73
HelloSuppressed	61 , 75

holddown timer	
RIP	119 , 121
HolddownTime	113 , 115
HoldDownTime	107

I

IfMtuIgnore	64 , 66
IfType	64 , 66
ImportAsExtern	69
in filter	19
InPolicy	111 , 113 , 115
Interface	111
interVRF redistribution	
configuring using EDM	129
IP address	
classes	12
in dotted-decimal notation	12
management port	23
multicast	12
subnet address	12
subnet mask	12
supernet	12
unicast	12
IP Address	57 , 61 , 62 , 75
IP routing	
BGPv4	11
brouter ports	16
OSPFv2	11
RIP	49
RIPv1	11
RIPv2	11
static route	17
VLAN	16
IpAddress	76

L

LastSpfRun	54
Listen	111 , 113 , 115
LSA	
areas and LSA types	42
LSDB	
viewing the LSDB using the ACLI	100
viewing using EDM	78
LSDB, external	
viewing using EDM	79
viewing using the ACLI	100
LsdbType	71
Lsid	78 , 79
LSRetransQLen	75

M

manual virtual link	41
Mask	71
MD5 key	92
Metric	62 , 64 , 66 , 70 , 76 , 80 , 117 , 129
MetricType	80 , 129

N

neighbor	19 , 27 , 28 , 32
Neighbor	73
nosend	110

O

Operation	107
Options	61 , 75
OriginateNewLsas	54
OSPF	
area border router	33
areas	29
AS	29
AS boundary router	33
AS boundary router (ASBR)	42
ASBR summary link advertisement	40
ASE link advertisement	40
ASE routes	41
autonomous system external (ASE) routes	41
backbone area	29
backup designated router	33
best path	44
broadcast interface	34
broadcast interfaces	34
configuring a host route using the ACLI	96
configuring for a port or VLAN using the ACLI	86
configuring global default metrics using EDM	56
configuring globally using EDM	54
configuring globally using the ACLI	83
configuring host routes using EDM	76
configuring on a brouter port using EDM	64
configuring on a VLAN using EDM	66
database description (DD) packets	39
designated router	33
enabling ASBR status using EDM	77
enabling automatic virtual links using EDM	73
enabling automatic virtual links using the ACLI	92
enabling globally using EDM	56

forcing shortest-path calculation updates using EDM	82	configuring ABR stub area metrics using EDM	70
forcing shortest-path calculation updates using the ACLI	104	configuring aggregate ranges using EDM	71
hello packets	39	configuring aggregate ranges using the ACLI	90
importing small stub routing domains	29	configuring manual virtual interface using EDM	73
interfaces description	34	configuring not-so-stubby areas using EDM	69
internal router	33	configuring on a VLAN or port using the ACLI	94
IP	39	configuring stubby areas using EDM	69
link-state acknowledgements	39	configuring using the ACLI	89
link-state advertisements	40	OSPF host route advertisements	
link-state database	27	on noncore areas	46
link-state request packets	39	OSPF interface	
link-state update packets	39	changing type using EDM	60
LSA	27	configuring metrics using EDM	62
manual aggregation	29	configuring NBMA neighbors using EDM	61
MD5	44	configuring using EDM	57
metrics	44	viewing using EDM	63
NBMA		OSPF neighbor	
forming adjacencies	34	configuring NBMA neighbors using the ACLI	97
poll interval	34	OSPF neighbors	
NBMA interface	34	managing using EDM	78
NBMA interfaces	34	OspfAction	54
neighbor	32	ospfNbmaNbrPermanence	61
network links advertisement	40	out filter	19
network summary link advertisement	40	OutPolicy	111 , 113 , 115
not so stubby area	29		
packet types	39	P	
passive interface	34	Poison	111 , 113 , 115
passive interfaces	34	PollInterval	57 , 64 , 66
redistribution	45	Priority	61
route advertisement	39	Protocol	80 , 117 , 129
router links advertisement	40		
router types	33	Q	
routing algorithm	28	Queries	107
security	44		
simple password	44	R	
specifying ASBRs	42	Receive	108
stub area	29	redistribution	18 , 19 , 44 , 45 , 49 , 80 , 101 , 103 , 117 , 124 , 125
transit areas	29	and CLIP	18
variable-length mask	39	configuring interVRF route redistribution for OSPF using the ACLI	103
viewing external LSDB using the ACLI	100	configuring interVRF route redistribution for RIP using the ACLI	125
viewing the external link-state database using EDM	79	configuring route redistribution to OSPF using EDM	80
viewing the link-state database using EDM	78		
viewing the LSDB using the ACLI	100		
viewing virtual neighbors using EDM	75		
virtual link	32 , 41		
virtual links	41		
OSPF area			
configuring a virtual interface using the ACLI	92		

configuring route redistribution to OSPF using the ACLI	101	accept	19
Configuring route redistribution to RIP using EDM ...	117	announce	19
configuring route redistribution to RIP using the ACLI	124	application order	19
OSPF	45	route preference	23
RIP	49	and alternative routes	23
redistribution filter	19	route preferences	23
reseller	131	RouteChanges	107
RetransInterval	57 , 73	RoutePolicy	80 , 117 , 129
Retransmission Queue Length	61	Router	61
retransmit interval	92	RouterId	79
retransmit interval, OSPF	86 , 94	RouterID	54 , 78
reverse path checking	23	RouteSource	80 , 117 , 129
Rfc1583Compatability	54	routing and security	23
RIP		RtrDeadInterval	57 , 64 , 66 , 73
advertise-when-down option	121	RtrId	75
advertisements	49	RtrPriority	57
configuring globally using EDM	107	RxNewLsas	54
configuring globally using the ACLI	119		
configuring on a brouter port using EDM	113	S	
configuring on a VLAN using EDM	115	Send	108
configuring on an interface using the ACLI	121	Sequence	78 , 79
distance vector protocol	49	showing OSPF port errors using the ACLI	88
forcing a RIP update using the ACLI	127	SpfHoldDownTime	54
hop count	49	SpfRuns	69
metric	49	SrcVrflD	80 , 117 , 129
multicast advertisements	49	State	61 , 73 , 75
poison reverse	121	static route	17
redistribution	49	black hole	17
routing table	49	High Availability mode	17
send modes	110	Status	62 , 70
split horizon	121	subnet mask	12
supply and listen settings	112	variable-length	12
TCP/IP route information	49	Subnets	80 , 129
UDP	49	supernet addressing	12
variable length subnet masks (VLSM)	49	Supply	111 , 113 , 115
RIP interface			
configuring using EDM	111	T	
RIP interface compatibility		TimeOutInterval	107 , 113 , 115
configuring using EDM	108	TOS	62 , 70 , 76
rip1	110	training	131
rip1comp	110	transit delay	92
rip2	110	transit delay, OSPF	86 , 94
RIPv1	49	TransitDelay	57 , 73
RIPv2	49	TrapEnable	54
RMON and CLIP	18	TriggeredUpdate	111
route acceptance policies		TriggeredUpdateEnable	113 , 115
applying using the ACLI	98	Type	57 , 78 , 79
route policy	19	Type 1 external metrics	44
		Type 2 external metrics	44

U

UDP

RIP	49
update timer, RIP	119
UpdateTime	107

V

VersionNumber	54
---------------------	--------------------

VLANs

routing between	16
-----------------------	--------------------

VLSM	12
------------	--------------------

VRF redistribution

configuring OSPF interVRF route redistribution	129
---	---------------------

configuring RIP interVRF route redistribution	129
---	---------------------

configuring using EDM	129
-----------------------------	---------------------

