Chapter 7
Configuring OSPF Version 3

Open Shortest Path First (OSPF) is a link-state routing protocol. OSPF uses link-state advertisements (LSAs) to update neighboring routers about its interfaces and information on those interfaces. The router floods LSAs to all neighboring routers to update them about the interfaces. Each router maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.

This chapter describes the following:

- The differences between OSPF versions 2 and 3.
- The link state advertisement types for OSPF version 3.
- How to configure OSPF version 3.
- How to display OSPF version 3 information and statistics.

OSPF Version 3

IPv6 supports OSPF version 3 (OSPFv3), which functions similarly to OSPF version 2, the current version that IPv4 supports, except for the following enhancements:

- Support for IPv6 addresses and prefixes.
- In general, you can configure several IPv6 addresses on a router interface. OSPFv3 imports all or none of the address prefixes configured on a router interface. You cannot select which addresses to import.
- You can run one instance of OSPF version 2 and one instance of OSPFv3 concurrently on a link.
- IPv6 link state advertisements (LSAs).

In addition, HP implements some new commands that are specific to OSPFv3. This section describes the commands that are specific to OSPFv3.

NOTE: Although OSPF versions 2 and 3 function similarly to each other, HP has implemented the user interface for each version independently of each other. Therefore, any configuration of OSPF version 2 features will not affect the configuration of OSPFv3 features and vice versa.

Link State Advertisement Types for OSPFv3

OSPFv3 supports the following types of LSAs:

- Router LSAs (Type 1)

- Network LSAs (Type 2)
- Interarea-prefix LSAs for ABRs (Type 3)
- Interarea-router LSAs for ASBRs (Type 4)
- Autonomous system external LSAs (Type 5)
- Link LSAs (Type 8)
- Intra-area prefix LSAs (Type 9)

For more information about these LSAs, see RFC 2740.

Configuring OSPFv3

To configure OSPFv3, you must do the following:

- Enable OSPFv3 globally.
- Assign OSPF areas.
- Assign router interfaces to an OSPF area.

The following configuration tasks are optional:

- Configure a virtual link between an ABR without a physical connection to a backbone area and the HP device in the same area with a physical connection to the backbone area.
- Change the reference bandwidth for the cost on OSPFv3 interfaces.
- Configure the redistribution of routes into OSPFv3.
- Configure default route origination.
- Modify the shortest path first (SPF) timers.
- Modify the administrative distances for OSPFv3 routes.
- Configure the OSPFv3 LSA pacing interval
- Modify how often the HP device checks on the elimination of the database overflow condition.
- Modify the external link state database limit.
- Modify the default values of OSPFv3 parameters for router interfaces.
- Disable or reenable OSPFv3 event logging.

Enabling OSPFv3

Before enabling the HP device to run OSPFv3, you must do the following:

- Enable the forwarding of IPv6 traffic on the HP device using the **ipv6 unicast-routing** command.
- Enable IPv6 on each interface over which you plan to enable OSPFv3. You enable IPv6 on an interface by configuring an IPv6 address or explicitly enabling IPv6 on that interface.

For more information about performing these configuration tasks, see “Configuring Basic IPv6 Connectivity” on page 3-1.

By default, OSPFv3 is disabled. To enable OSPFv3, you must enable it globally.

To enable OSPFv3 globally, enter the following command:

```
ProCurveRS(config-ospf-router)#ipv6 router ospf
ProCurveRS(config-ospf6-router)#
```

After you enter this command, the HP device enters the IPv6 OSPF configuration level, where you can access several commands that allow you to configure OSPFv3.

Syntax: [no] ipv6 router ospf

To disable OSPFv3, enter the **no** form of this command. If you disable OSPFv3, the HP device removes all the configuration information for the disabled protocol from the running-config. Moreover, when you save the configuration to the startup-config file after disabling one of these protocols, all the configuration information for the disabled protocol is removed from the startup-config file.

The CLI displays a warning message such as the following:

```
ProCurveRS(config-ospf6-router)# no ipv6 router ospf
ipv6 router ospf mode now disabled. All ospf config data will be lost when writing
to flash!
```

If you have disabled the protocol but have not yet saved the configuration to the startup-config file and reloaded the software, you can restore the configuration information by re-entering the command to enable the protocol (for example, **ipv6 router ospf**). If you have already saved the configuration to the startup-config file and reloaded the software, the information is gone. If you are testing an OSPF configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup-config file containing the protocol's configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

Assigning OSPFv3 Areas

After OSPFv3 is enabled, you can assign OSPFv3 areas. You can assign an IPv4 address or a number as the ***area ID*** for each area. The area ID is representative of all IPv6 addresses (subnets) on a router interface. Each router interface can support one area.

An area can be **normal** or a **stub**.

- Normal – OSPF routers within a normal area can send and receive External Link State Advertisements (LSAs).
- Stub – OSPF routers within a stub area cannot send or receive External LSAs. In addition, OSPF routers in a stub area must use a default route to the area's Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) to send traffic out of the area.

For example, to set up OSPFv3 areas 0.0.0.0, 200.5.0.0, 192.5.1.0, and 195.5.0.0, enter the following commands:

```
ProCurveRS(config-ospf6-router)# area 0.0.0.0
ProCurveRS(config-ospf6-router)# area 200.5.0.0
ProCurveRS(config-ospf6-router)# area 192.5.1.0
ProCurveRS(config-ospf6-router)# area 195.5.0.0
```

Syntax: [no] area <number> | <ipv4-address>

The <number> | <ipv4-address> parameter specifies the area number, which can be a number or in IPv4 address format. If you specify a number, the number can be from 0 – 2,147,483,647.

NOTE: You can assign one area on a router interface.

Assigning a Totally Stubby Area

By default, the HP device sends summary LSAs (LSA type 3) into stub areas. You can further reduce the number of LSAs sent into a stub area by configuring the HP device to stop sending summary LSAs into the area. You can disable the summary LSAs when you are configuring the stub area or later after you have configured the area.

This feature disables origination of summary LSAs into a stub area, but the HP device still accepts summary LSAs from OSPF neighbors and floods them to other areas. The HP device can form adjacencies with other routers regardless of whether summarization is enabled or disabled for areas on each router.

When you disable the summary LSAs, the change takes effect immediately. If you apply the option to a previously configured area, the router flushes all of the summary LSAs it has generated (as an ABR) from the area.

NOTE: This feature applies only when the HP device is configured as an Area Border Router (ABR) for the area. To completely prevent summary LSAs from being sent to the area, disable the summary LSAs on each OSPF router that is an ABR for the area.

For example, to disable summary LSAs for stub area 40 and specify an additional metric of 99, enter the following command:

```
ProCurveRS(config-ospf6-router)# area 40 stub 99 no-summary
```

Syntax: area <number> | <ipv4-address> stub <metric> [no-summary]

The <number> | <ipv4-address> parameter specifies the area number, which can be a number or in IPv4 address format. If you specify a number, the number can be from 0 – 2,147,483,647.

The **stub** <metric> parameter specifies an additional cost for using a route to or from this area and can be from 1 – 16777215. There is no default. Normal areas do not use the cost parameter.

The **no-summary** parameter applies only to stub areas and disables summary LSAs from being sent into the area.

Assigning Interfaces to an Area

After you define OSPFv3 areas, you must assign router interfaces to the areas. All router interfaces must be assigned to one of the defined areas on an OSPF router. When an interface is assigned to an area, all corresponding subnets on that interface are automatically included in the assignment.

For example, to assign Ethernet interface 3/1 to area 192.5.0.0, enter the following commands:

```
ProCurveRS(config)# interface Ethernet 3/1
ProCurveRS(config-if-e100-3/1)# ipv6 ospf area 195.5.0.0
```

Syntax: [no] ipv6 ospf area <number> | <ipv4-address>

The <number> | <ipv4-address> parameter specifies the area number, which can be a number or in IPv4 address format. If you specify a number, the number can be from 0 – 2,147,483,647.

To remove the interface from the specified area, use the **no** form of this command.

Configuring Virtual Links

All ABRs must have either a direct or indirect link to an OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to a backbone area, you can configure a virtual link from the ABR to another router within the same area that has a physical connection to the backbone area.

The path for a virtual link is through an area shared by the neighbor ABR (router with a physical backbone connection) and the ABR requiring a logical connection to the backbone.

Two parameters must be defined for all virtual links—transit area ID and neighbor router.

- The transit area ID represents the shared area of the two ABRs and serves as the connection point between the two routers. This number should match the area ID value.
- When assigned from the router interface requiring a logical connection, the neighbor router field is the router ID (IPv4 address) of the router that is physically connected to the backbone. When assigned from the router interface with the physical connection, the neighbor router is the router ID (IPv4) address of the router requiring a logical connection to the backbone.

NOTE: By default, the HP router ID is the IPv4 address configured on the lowest numbered loopback interface. If the HP device does not have a loopback interface, the default router ID is the lowest numbered IPv4 address configured on the device.

NOTE: When you establish an area virtual link, you must configure it on both of the routers (both ends of the virtual link).

For example, imagine that ABR1 in areas 1 and 2 is cut off from the backbone area (area 0). To provide backbone access to ABR1, you can add a virtual link between ABR1 and ABR2 in area 1 using area 1 as a transit area. To configure the virtual link, you define the link on the router that is at each end of the link. No configuration for the virtual link is required on the routers in the transit area.

To define the virtual link on ABR1, enter the following command on ABR1:

```
ProCurveRS(config-ospf6-router)# area 1 virtual-link 209.157.22.1
```

To define the virtual link on ABR2, enter the following command on ABR2:

```
ProCurveRS(config-ospf6-router)# area 1 virtual-link 10.0.0.1
```

Syntax: area <number> | <ipv4-address> virtual-link <router-id>

The **area** <number> | <ipv4-address> parameter specifies the transit area.

The <router-id> parameter specifies the router ID of the OSPF router at the remote end of the virtual link. To display the router ID on a router, enter the **show ip** command.

Assigning a Virtual Link Source Address

When routers at both ends of a virtual link need to communicate with one another, the source address included in the packets must be a global IPv6 address. Therefore, you must determine the global IPv6 address to be used by the routers for communication across the virtual link. You can specify that a router uses the IPv6 global address assigned to one of its interfaces.

For example, to specify the global IPv6 address assigned to Ethernet interface 3/1 on ABR1 as the source address for the virtual link on ABR1, enter the following command on ABR1:

```
ProCurveRS(config-ospf6-router)# virtual-link-if-address interface ethernet 3/1
```

To specify the global IPv6 address assigned to tunnel interface 1 on ABR2 as the source address for the virtual link on ABR2, enter the following command on ABR2:

```
ProCurveRS(config-ospf6-router)# virtual-link-if-address interface tunnel 1
```

Syntax: virtual-link-if-address interface ethernet <port> | loopback <number> | tunnel <number> | ve <number>

The **ethernet | loopback | tunnel | ve** parameter specifies the interface from which the router derives the source IPv6 address for communication across the virtual link. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a loopback, tunnel, or VE interface, also specify the number associated with the respective interface.

To delete the source address for the virtual link, use the **no** form of this command.

Modifying Virtual Link Parameters

You can modify the following virtual link parameters:

- Dead-interval: The number of seconds that a neighbor router waits for a hello packet from the current router before declaring the router is down. The range is 1 – 65535 seconds. The default is 40 seconds.
- Hello-interval: The length of time between the transmission of hello packets. The range is 1 – 65535 seconds. The default is 10 seconds.
- Retransmit-interval: The interval between the re-transmission of link state advertisements to router adjacencies for this interface. The range is 0 – 3600 seconds. The default is 5 seconds.
- Transmit-delay: The period of time it takes to transmit Link State Update packets on the interface. The range is 0 – 3600 seconds. The default is 1 second.

NOTE: The values of the **dead-interval** and **hello-interval** parameters must be the same at both ends of a virtual link. Therefore, if you modify the values of these parameters at one end of a virtual link, you must remember to make the same modifications on the other end of the link.

The values of the other virtual link parameters do not require synchronization.

For example, to change the dead interval to 60 seconds on the virtual links defined on ABR1 and ABR2, enter the following command on ABR1:

```
ProCurveRS(config-ospf6-router)# area 1 virtual-link 209.157.22.1
dead-interval 60
```

Enter the following command on ABR2:

```
ProCurveRS(config-ospf6-router)# area 1 virtual-link 10.0.0.1 dead-interval 60
```

Syntax: area <number> | <ipv4-address> virtual-link <router-id> [dead-interval <seconds> | hello-interval <seconds> | retransmit-interval <seconds> | transmit-delay <seconds>]

The **area** <number> | <ipv4-address> parameter specifies the transit area.

The <router-id> parameter specifies the router ID of the OSPF router at the remote end of the virtual link. To display the router ID on a router, enter the **show ip** command.

The **dead-interval**, **hello-interval**, **retransmit-interval**, and **transmit-delay** parameters are discussed earlier in this section.

Changing the Reference Bandwidth for the Cost on OSPFv3 Interfaces

Each interface on which OSPFv3 is enabled has a cost associated with it. The HP device advertises its interfaces and their costs to OSPFv3 neighbors. For example, if an interface has an OSPF cost of ten, the HP device advertises the interface with a cost of ten to other OSPF routers.

By default, an interface's OSPF cost is based on the port speed of the interface. The software uses the following formula to calculate the cost:

Cost = reference-bandwidth/interface-speed

By default, the reference bandwidth is 100 Mbps. If the resulting cost is less than 1, the software rounds the cost up to 1. The default reference bandwidth results in the following costs:

- 10 Mbps port's cost = 100/10 = 10
- 100 Mbps port's cost = 100/100 = 1
- 1000 Mbps port's cost = 100/1000 = 0.10, which is rounded up to 1
- 155 Mbps port's cost = 100/155 = 0.65, which is rounded up to 1
- 622 Mbps port's cost = 100/622 = 0.16, which is rounded up to 1
- 2488 Mbps port's cost = 100/2488 = 0.04, which is rounded up to 1

The bandwidth for interfaces that consist of more than one physical port is calculated as follows:

- Trunk group – The combined bandwidth of all the ports.
- Virtual (Ethernet) interface – The combined bandwidth of all the ports in the port-based VLAN that contains the virtual interface.

You can change the default reference bandwidth from 100 Mbps to a value from 1 – 4294967 Mbps.

If a change to the reference bandwidth results in a cost change to an interface, the HP device sends a link state update to update the costs of interfaces advertised by the HP device.

NOTE: If you specify the cost for an individual interface, the cost you specify overrides the cost calculated by the software.

Some interface types are not affected by the reference bandwidth and always have the same cost regardless of the reference bandwidth in use:

- The cost of a loopback interface is always 0.
- The cost of a virtual link is calculated using the Shortest Path First (SPF) algorithm and is not affected by the auto-cost feature.

- The bandwidth for tunnel interfaces is 9 Kbps and is not affected by the auto-cost feature.

For example, to change the reference bandwidth to 500, enter the following command:

```
ProCurveRS(config-ospf6-router)# auto-cost reference-bandwidth 500
```

The reference bandwidth specified in this example results in the following costs:

- 10 Mbps port's cost = 500/10 = 50
- 100 Mbps port's cost = 500/100 = 5
- 1000 Mbps port's cost = 500/1000 = 0.5, which is rounded up to 1
- 155 Mbps port's cost = 500/155 = 3.23, which is rounded up to 4
- 622 Mbps port's cost = 500/622 = 0.80, which is rounded up to 1
- 2488 Mbps port's cost = 500/2488 = 0.20, which is rounded up to 1

The costs for 10 Mbps, 100 Mbps, and 155 Mbps ports change as a result of the changed reference bandwidth. Costs for higher-speed interfaces remain the same.

Syntax: [no] auto-cost reference-bandwidth <number>

The <number> parameter specifies the reference bandwidth and can be a value from 1 – 4294967. The default is 100.

To restore the reference bandwidth to its default value and thus restore the default costs of interfaces to their default values, enter the **no** form of this command.

Redistributing Routes into OSPFv3

In addition to specifying which routes are redistributed into OSPFv3, you can configure the following aspects related to route redistribution:

- Default metric.
- Metric type.
- Advertisement of an external aggregate route.

Configuring Route Redistribution into OSPFv3

You can configure the HP device to redistribute routes from the following sources into OSPFv3:

- IPv6 static routes.
- Directly connected IPv6 networks.
- BGP4+.
- RIPng.

You can redistribute routes in the following ways:

- By route types, for example, the HP device redistributes all IPv6 static and RIPng routes.
- By using a route map to filter which routes to redistribute, for example, the HP device redistributes specified IPv6 static and RIPng routes only.

For example, to configure the redistribution of all IPv6 static and RIPng, enter the following commands:

```
ProCurveRS(config-ospf6-router)# redistribute static
ProCurveRS(config-ospf6-router)# redistribute rip
```

Syntax: [no] redistribute bgp | connected | rip | static [metric <number> | metric-type <type>]

The **bgp** | **connected** | **rip** | **static** keywords specify the route source.

The **metric** <number> parameter specifies the metric used for the redistributed route. If a value is not specified for this option, and the value for the **default-metric** command is set to 0, its default metric, then routes redistributed

from the various routing protocols will have the metric value of the protocol from which they are redistributed. For information about the **default-metric** command, see “Modifying Default Metric for Routes Redistributed into OSPF Version 3” on page 7-9.

The **metric-type** <type> parameter specifies an OSPF metric type for the redistributed route. You can specify external type 1 or external type 2. If a value is not specified for this option, the HP device uses the value specified by the **metric-type** command. For information about modifying the default metric type using the **metric-type** command, see “Modifying Metric Type for Routes Redistributed into OSPF Version 3” on page 7-9.

For example, to configure a route map and use it for redistribution of routes into OSPFv3, enter commands such as the following:

```
ProCurveRS(config)# ipv6 route 2001:1::/32 4823:eoff:343e::23
ProCurveRS(config)# ipv6 route 2001:2::/32 4823:eoff:343e::23
ProCurveRS(config)# ipv6 route 2001:3::/32 4823:eoff:343e::23 metric 5
ProCurveRS(config)# route-map abc permit 1
ProCurveRS(config-routemap abc)# match metric 5
ProCurveRS(config-routemap abc)# set metric 8
ProCurveRS(config-routemap abc)# ipv6 router ospf
ProCurveRS(config-ospf6-router)# redistribute static route-map abc
```

The commands in this example configure some static IPv6 routes and a route map, and use the route map for redistributing the static IPv6 routes into OSPFv3.

The **ipv6 route** commands configure the static IPv6 routes. The **route-map** command begins configuration of a route map called “abc”. The number indicates the route map entry (called the “instance”) you are configuring. A route map can contain multiple entries. The software compares packets to the route map entries in ascending numerical order and stops the comparison once a match is found.

The **match** command in the route map matches on routes that have 5 for their metric value (cost). The **set** command changes the metric in routes that match the route map to 8.

The **redistribute** command configures the redistribution of static IPv6 routes into OSPFv3, and uses route map “abc“ to control the routes that are redistributed. In this example, the route map allows a static IPv6 route to be redistributed into OSPF only if the route has a metric of 5, and changes the metric to 8 before placing the route into the OSPF route redistribution table.

Syntax: [no] redistribute bgp | connected | rip | static [route-map <map-name>]

The **bgp** | **connected** | **rip** | **static** keywords specify the route source.

The **route-map** <map-name> parameter specifies the route map name. The following match parameters are valid for OSPFv3 redistribution:

- **match ip address** | **next-hop** <acl-number>
- **match metric** <number>
- **match tag** <tag-value>

The following set parameters are valid for OSPF redistribution:

- **set ip next hop** <ipv4-address>
- **set metric** [+ | -] <number> | none
- **set metric-type** type-1 | type-2
- **set tag** <tag-value>

NOTE: You must configure the route map before you configure a redistribution filter that uses the route map.

NOTE: When you use a route map for route redistribution, the software disregards the permit or deny action of the route map.

NOTE: For an external route that is redistributed into OSPFv3 through a route map, the metric value of the route remains the same unless the metric is set by a **set metric** command inside the route map or the **default-metric** <num> command. For a route redistributed without using a route map, the metric is set by the metric parameter if set or the **default-metric** <num> command if the metric parameter is not set.

Modifying Default Metric for Routes Redistributed into OSPF Version 3

The default metric is a global parameter that specifies the cost applied by default to routes redistributed into OSPFv3. The default value is 0.

If the **metric** parameter for the **redistribute** command is not set and the **default-metric** command is set to 0, its default value, then routes redistributed from the various routing protocols will have the metric value of the protocol from which they are redistributed. For information about the **redistribute** command, see “Configuring Route Redistribution into OSPFv3” on page 7-7.

NOTE: You also can define the cost on individual interfaces. The interface cost overrides the default cost. For information about defining the cost on individual interfaces, see “Modifying OSPFv3 Interface Defaults” on page 7-16 and “Changing the Reference Bandwidth for the Cost on OSPFv3 Interfaces” on page 7-6.

To assign a default metric of 4 to all routes imported into OSPFv3, enter the following command:

```
ProCurveRS(config-ospf6-router)# default-metric 4
```

Syntax: [no] default-metric <number>

You can specify a value from 0 – 65535. The default is 0.

To restore the default metric to the default value, use the **no** form of this command.

Modifying Metric Type for Routes Redistributed into OSPF Version 3

The HP device uses the **metric-type** parameter by default for all routes redistributed into OSPFv3 unless you specify a different metric type for individual routes using the **redistribute** command. (For more information about using the **redistribute** command, see “Redistributing Routes into OSPFv3” on page 7-7.)

A type 1 route specifies a small metric (two bytes), while a type 2 route specifies a big metric (three bytes). The default value is type 2.

To modify the default value of type 2 to type 1, enter the following command:

```
ProCurveRS(config-ospf6-router)# metric-type type1
```

Syntax: [no] metric-type type1 | type2

To restore the metric type to the default value, use the **no** form of this command.

Configuring External Route Summarization

When the HP device is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to advertise one external route as an aggregate for all redistributed routes that are covered by a specified IPv6 address range.

When you configure an address range, the range takes effect immediately. All the imported routes are summarized according to the configured address range. Imported routes that have already been advertised and that fall within the range are flushed out of the AS and a single route corresponding to the range is advertised.

If a route that falls within a configured address range is imported by the HP device, no action is taken if the device has already advertised the aggregate route; otherwise, the device advertises the aggregate route. If an imported route that falls within a configured address range is removed by the device, no action is taken if there are other imported route(s) that fall with in the same address range; otherwise the aggregate route is flushed.

You can configure up to 32 address ranges. The HP device sets the forwarding address of the aggregate route to zero and sets the tag to zero.

If you delete an address range, the advertised aggregate route is flushed and all imported routes that fall within the range are advertised individually.

If an external link state database overflow (LSDB) condition occurs, all aggregate routes are flushed out of the AS, along with other external routes. When the HP device exits the external LSDB overflow condition, all the imported routes are summarized according to the configured address ranges.

NOTE: If you use redistribution filters in addition to address ranges, the HP device applies the redistribution filters to routes first, then applies them to the address ranges.

NOTE: If you disable redistribution, all the aggregate routes are flushed, along with other imported routes.

NOTE: This option affects only imported, type 5 external routes. A single type 5 LSA is generated and flooded throughout the AS for multiple external routes.

To configure the summary address 2201::/24 for routes redistributed into OSPFv3, enter the following command:

```
ProCurveRS(config-ospf6-router)# summary-address 2201::/24
```

In this example, the summary prefix 2201::/24 includes addresses 2201::/1 through 2201::/24. Only the address FEC0::/24 is advertised in an external link-state advertisement.

Syntax: summary-address <ipv6-prefix>/<prefix-length>

You must specify the <ipv6-prefix> parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

You must specify the <prefix-length> parameter as a decimal value. A slash mark (/) must follow the <ipv6-prefix> parameter and precede the <prefix-length> parameter.

Filtering OSPFv3 Routes

You can filter the routes to be placed in the OSPFv3 route table by configuring distribution lists. OSPFv3 distribution lists can be applied globally or to an interface.

The functionality of OSPFv3 distribution lists is similar to that of OSPFv2 distribution lists. However, unlike OSPFv2 distribution lists, which filter routes based on criteria specified in an Access Control List (ACL), OSPFv3 distribution lists can filter routes using information specified in an IPv6 prefix list or a route map.

Configuration Examples

The following sections show examples of filtering OSPFv3 routes using prefix lists globally and for a specific interface, as well as filtering OSPFv3 routes using a route map.

You can configure the device to use all three types of filtering. When you do this, filtering using route maps has higher priority over filtering using global prefix lists. Filtering using prefix lists for a specific interface has lower priority than the other two filtering methods.

The example in this section assume the following routes are in the OSPFv3 route table:

```
ProCurveRS# show ipv6 ospf route

  Current Route count: 5
    Intra: 3 Inter: 0 External: 2 (Type1 0/Type2 2)
    Equal-cost multi-path: 0
    Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 3001::/64                      --------- 0.0.0.1            0 0
    ::                             ve 10
*E2 3010::/64                      --------- 0.0.0.0           10 0
    fe80::2e0:52ff:fe00:10         ve 10
*IA 3015::/64                      V6E---R-- 0.0.0.0           11 0
    fe80::2e0:52ff:fe00:10         ve 10
*IA 3020::/64                      --------- 0.0.0.0           10 0
    ::                             ve 11
*E2 6001:5000::/64                 --------- 0.0.0.0           10 0
    fe80::2e0:52ff:fe00:10         ve 10
```

Configuring an OSPFv3 Distribution List Using an IPv6 Prefix List as Input

The following example illustrates how to use an IPv6 prefix list is used to filter OSPFv3 routes.

To specify an IPv6 prefix list called filterOspfRoutes that denies route 3010::/64, enter the following commands:

```
ProCurveRS(config)# ipv6 prefix-list  filterOspfRoutes seq 5 deny 3010::/64
ProCurveRS(config)# ipv6 prefix-list  filterOspfRoutes seq 7 permit ::/0 ge 1 le 128
```

Syntax: ipv6 prefix-list <name> [seq <seq-value>] [description <string>] deny | permit <ipv6-addr>/<mask-bits> [ge <ge-value>] [le <le-value>]

See "Defining IP Prefix Lists" in the *Advanced Configuration and Management Guide for ProCurve 9300/9400 Series Routing Switches* for information on configuring prefix lists.

To configure a distribution list that applies the filterOspfRoutes prefix list globally:

```
ProCurveRS(config)# ipv6 router ospf
ProCurveRS(config-ospf6-router)# distribute-list prefix-list filterOspfRoutes in
```

Syntax: [no] distribute-list prefix-list <name> in [<interface>]

After this distribution list is configured, route 3010::/64 would be omitted from the OSPFv3 route table:

```
ProCurveRS# show ipv6 ospf route

  Current Route count: 4
    Intra: 3 Inter: 0 External: 1 (Type1 0/Type2 1)
    Equal-cost multi-path: 0
    Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 3001::/64                      --------- 0.0.0.1            0 0
    ::                             ve 10
*IA 3015::/64                      V6E---R-- 0.0.0.0           11 0
    fe80::2e0:52ff:fe00:10         ve 10
*IA 3020::/64                      --------- 0.0.0.0           10 0
    ::                             ve 11
*E2 6001:5000::/64                 --------- 0.0.0.0           10 0
    fe80::2e0:52ff:fe00:10         ve 10
```

The following commands specify an IPv6 prefix list called filterOspfRoutesVe that denies route 3015::/64:

```
ProCurveRS(config)# ipv6 prefix-list filterOspfRoutesVe seq 5 deny 3015::/64
ProCurveRS(config)# ipv6 prefix-list filterOspfRoutesVe seq 10 permit ::/0 ge 1 le
128
```

The following commands configure a distribution list that applies the filterOspfRoutesVe prefix list to routes pointing to virtual interface 10:

```
ProCurveRS(config)# ipv6 router ospf
ProCurveRS(config-ospf6-router)# distribute-list prefix-list filterOspfRoutes in ve
10
```

After this distribution list is configured, route 3015::/64, pointing to virtual interface 10, would be omitted from the OSPFv3 route table:

```
ProCurveRS# show ipv6 ospf route

  Current Route count: 4
    Intra: 3 Inter: 0 External: 1 (Type1 0/Type2 1)
    Equal-cost multi-path: 0
    Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 3001::/64                      --------- 0.0.0.1            0 0
    ::                             ve 10
*E2 3010::/64                      --------- 0.0.0.0           10 0
    fe80::2e0:52ff:fe00:10         ve 10
*IA 3020::/64                      --------- 0.0.0.0           10 0
    ::                             ve 11
*E2 6001:5000::/64                 --------- 0.0.0.0           10 0
    fe80::2e0:52ff:fe00:10         ve 10
```

Configuring an OSPFv3 Distribution List Using a Route Map as Input

The following commands configure a route map that matches internal routes:

```
ProCurveRS(config)# route-map allowInternalRoutes permit 10
ProCurveRS(config-routemap allowInternalRoutes)# match route-type internal
```

See "Hardware-Based Policy-Based Routing (PBR)" in the *Advanced Configuration and Management Guide for ProCurve 9300/9400 Series Routing Switches* for information on configuring route maps.The following commands configure a distribution list that applies the allowInternalRoutes route map globally to OSPFv3 routes:

```
ProCurveRS(config)# ipv6 router ospf
ProCurveRS(config-ospf6-router)# distribute-list route-map allowinternalroutes in
```

Syntax: [no] distribute-list route-map <name> in

After this distribution list is configured, the internal routes would be included, and the external routes would be omitted from the OSPFv3 route table:

```
ProCurveRS# show ipv6 ospf route

  Current Route count: 3
    Intra: 3 Inter: 0 External: 0 (Type1 0/Type2 0)
    Equal-cost multi-path: 0
    Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 3001::/64                      --------- 0.0.0.1           0 0
    ::                             ve 10
*IA 3015::/64                      V6E---R-- 0.0.0.0          11 0
    fe80::2e0:52ff:fe00:10         ve 10
*IA 3020::/64                      --------- 0.0.0.0          10 0
    ::                             ve 11
```

Configuring Default Route Origination

When the HP device is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to automatically generate a default external route into an OSPFv3 routing domain. This feature is called “default route origination” or “default information origination.”

By default, the HP device does not advertise the default route into the OSPFv3 domain. If you want the device to advertise the OSPF default route, you must explicitly enable default route origination.

When you enable OSPF default route origination, the device advertises a type 5 default route that is flooded throughout the AS (except stub areas).

The device advertises the default route into OSPF even if OSPF route redistribution is not enabled, and even if the default route is learned through an IBGP neighbor.

NOTE: The HP device does not advertise the OSPF default route, regardless of other configuration parameters, unless you explicitly enable default route origination.

If default route origination is enabled and you disable it, the default route originated by the device is flushed. Default routes generated by other OSPF routers are not affected. If you re-enable the feature, the feature takes effect immediately and thus does not require you to reload the software.

For example, to create and advertise a default route with a metric of 2 and as a type 1 external route, enter the following command:

```
ProCurveRS(config-ospf6-router)# default-information-originate always metric 2
metric-type type1
```

Syntax: [no] default-information-originate [always] [metric <value>] [metric-type <type>]

The **always** keyword originates a default route regardless of whether the device has learned a default route. This option is disabled by default.

The **metric** <value> parameter specifies a metric for the default route. If this option is not used, the value of the **default-metric** command is used for the route. For information about this command, see “Modifying Default Metric for Routes Redistributed into OSPF Version 3” on page 7-9.

The **metric-type** <type> parameter specifies the external link type associated with the default route advertised into the OSPF routing domain. The <type> can be one of the following:

- 1 – Type 1 external route
- 2 – Type 2 external route

If you do not use this option, the default redistribution metric type is used for the route type.

NOTE: If you specify a metric and metric type, the values you specify are used even if you do not use the always option.

To disable default route origination, enter the **no** form of the command.

Modifying Shortest Path First Timers

The HP device uses the following timers when calculating the shortest path for OSPFv3 routes:

- SPF delay – When the HP device receives a topology change, the software waits before it starts a Shortest Path First (SPF) calculation. By default, the software waits 5 seconds. You can configure the SPF delay to a value from 0 – 65535 seconds. If you set the SPF delay to 0 seconds, the software immediately begins the SPF calculation after receiving a topology change.
- SPF hold time – The HP device waits a specific amount of time between consecutive SPF calculations. By default, the device waits 10 seconds. You can configure the SPF hold time to a value from 0 – 65535 seconds. If you set the SPF hold time to 0 seconds, the software does not wait between consecutive SPF calculations.

You can set the SPF delay and hold time to lower values to cause the device to change to alternate paths more quickly if a route fails. Note that lower values for these parameters require more CPU processing time.

You can change one or both of the timers.

NOTE: If you want to change only one of the timers, for example, the SPF delay timer, you must specify the new value for this timer as well as the current value of the SPF hold timer, which you want to retain. The HP device does not accept only one timer value.

To change the SPF delay to 10 seconds and the SPF hold to 20 seconds, enter the following command:

```
ProCurveRS(config-ospf6-router)# timers spf 10 20
```

Syntax: timers spf <delay> <hold-time>

For the <delay> and <hold-time> parameters, specify a value from 0 – 65535 seconds.

To set the timers back to their default values, enter the **no** version of this command.

Modifying Administrative Distance

The HP device can learn about networks from various protocols, including BGP4+ and OSPFv3. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. By default, the administrative distance for OSPFv3 routes is 110.

The device selects one route over another based on the source of the route information. To do so, the device can use the administrative distances assigned to the sources. You can influence the device's decision by changing the default administrative distance for OSPFv3 routes.

Configuring Administrative Distance Based on Route Type

You can configure a unique administrative distance for each type of OSPFv3 route. For example, you can use this feature to influence the HP device to prefer a static route over an OSPF inter-area route and to prefer OSPF intra-area routes to static routes.

The distance you specify influences the choice of routes when the device has multiple routes to the same network from different protocols. The device prefers the route with the lower administrative distance.

You can specify unique default administrative distances for the following OSPFv3 route types:

- Intra-area routes
- Inter-area routes
- External routes

The default for all of these OSPFv3 route types is 110.

NOTE: This feature does not influence the choice of routes within OSPFv3. For example, an OSPF intra-area route is always preferred over an OSPF inter-area route, even if the intra-area route's distance is greater than the inter-area route's distance.

For example, to change the default administrative distances for intra-area routes to 80, inter-area routes to 90, and external routes to 100, enter the following commands:

```
ProCurveRS(config-ospf6-router)# distance intra-area 80
ProCurveRS(config-ospf6-router)# distance inter-area 90
ProCurveRS(config-ospf6-router)# distance external 100
```

Syntax: distance external | inter-area | intra-area <distance>

The **external** | **inter-area** | **intra-area** keywords specify the route type for which you are changing the default administrative distance.

The <distance> parameter specifies the new distance for the specified route type. You can specify a value from 1 – 255.

To reset the administrative distance of a route type to its system default, enter the **no** form of this command.

Configuring the OSPFv3 LSA Pacing Interval

The HP device paces OSPFv3 LSA refreshes by delaying the refreshes for a specified time interval instead of performing a refresh each time an individual LSA's refresh timer expires. The accumulated LSAs constitute a group, which the HP device refreshes and sends out together in one or more packets.

The pacing interval, which is the interval at which the HP device refreshes an accumulated group of LSAs, is configurable to a range from 10 – 1800 seconds (30 minutes). The default is 240 seconds (four minutes). Thus, every four minutes, the HP device refreshes the group of accumulated LSAs and sends the group together in the same packet(s).

The pacing interval is inversely proportional to the number of LSAs the HP device is refreshing and aging. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval enhances performance. If you have a very small database (40 – 100 LSAs), increasing the pacing interval to 10 – 20 minutes might enhance performance only slightly.

To change the OSPFv3 LSA pacing interval to two minutes (120 seconds), enter the following command:

```
ProCurveRS(config)# ipv6 router ospf
ProCurveRS(config-ospf6-router)# timers lsa-group-pacing 120
```

Syntax: [no] timers lsa-group-pacing <seconds>

The <seconds> parameter specifies the number of seconds and can be from 10 – 1800 (30 minutes). The default is 240 seconds (four minutes).

To restore the pacing interval to its default value, use the **no** form of the command:

Modifying Exit Overflow Interval

If a database overflow condition occurs on the HP device, the device eliminates the condition by removing entries that originated on the device. The exit overflow interval allows you to set how often a device checks to see if the overflow condition has been eliminated. The default value is 0. If the configured value of the database overflow interval is 0, then the device never leaves the database overflow condition.

For example, to modify the exit overflow interval to 60 seconds, enter the following command:

```
ProCurveRS(config-ospf6-router)# database-overflow-interval 60
```

Syntax: database-overflow-interval <seconds>

The <seconds> parameter can be a value from 0 – 86400 seconds (24 hours).

To reset the exit overflow interval to its system default, enter the **no** form of this command.

Modifying External Link State Database Limit

By default, the link state database can hold a maximum of 2000 entries for external (type 5) LSAs. You can change the maximum number of entries from 500 – 8000. After changing this limit, make sure to save the running-config file and reload the software. The change does not take effect until you reload or reboot the software.

For example, to change the maximum number entries from the default of 2000 to 3000, enter the following command:

```
ProCurveRS(config-ospf6-router)# external-lsdb-limit 3000
```

Syntax: external-lsdb-limit <entries>

The <entries> parameter can be a numerical value from 500 – 8000 seconds.

To reset the maximum number of entries to its system default, enter the **no** form of this command.

Modifying OSPFv3 Interface Defaults

OSPFv3 has interface parameters that you can configure. For simplicity, each of these parameters has a default value. No change to these default values is required except as needed for specific network configurations.

You can modify the default values for the following OSPF interface parameters:

- Cost: Indicates the overhead required to send a packet across an interface. You can modify the cost to differentiate between 100 Mbps and 1000 Mbps (1 Gbps) links. The command syntax is **ipv6 ospf cost** <number>. The default cost is calculated by dividing 100 million by the bandwidth. For 10 Mbps links, the cost is 10. The cost for both 100 Mbps and 1000 Mbps links is 1, because the speed of 1000 Mbps was not in use at the time the OSPF cost formula was devised.
- Dead-interval: Indicates the number of seconds that a neighbor router waits for a hello packet from the current router before declaring the router down. The command syntax is **ipv6 ospf dead-interval** <seconds>. The value can be from 1 – 2147483647 seconds. The default is 40 seconds.
- Hello-interval: Represents the length of time between the transmission of hello packets. The command syntax is **ipv6 ospf hello-interval** <seconds>. The value can be from 1 – 65535 seconds. The default is 10 seconds.
- Instance: Indicates the number of OSPFv3 instances running on an interface. The command syntax is **ipv6 ospf instance** <number>. The value can be from 0 – 255. The default is 1.
- MTU-ignore: Allows you to disable a check that verifies the same MTU is used on an interface shared by neighbors. The command syntax is **ipv6 ospf mtu-ignore**. By default, the mismatch detection is enabled.
- Network: Allows you to configure the OSPF network type. The command syntax is **ipv6 ospf network** [**point-to-multipoint**]. The default setting of the parameter depends on the network type.
- Passive: When you configure an OSPF interface to be passive, that interface does not send or receive OSPF route updates. This option affects all IPv6 subnets configured on the interface. The command syntax is **ipv6 ospf passive**. By default, all OSPF interfaces are active and thus can send and receive OSPF route information. Since a passive interface does not send or receive route information, the interface is in effect a stub network.
- Priority: Allows you to modify the priority of an OSPF router. The priority is used when selecting the designated router (DR) and backup designated routers (BDRs). The command syntax is **ipv6 ospf priority** <number>. The value can be from 0 – 255. The default is 1. If you set the priority to 0, the router does not participate in DR and BDR election.
- Retransmit-interval: The time between retransmissions of LSAs to adjacent routers for an interface. The command syntax is **ipv6 ospf retransmit-interval** <seconds>. The value can be from 0 – 3600 seconds. The default is 5 seconds.
- Transmit-delay: The time it takes to transmit Link State Update packets on this interface. The command syntax is **ipv6 ospf transmit-delay** <seconds>. The value can be from 0 – 3600 seconds. The default is 1 second.

Disabling or Reenabling Event Logging

OSPFv3 does not currently support the generation of SNMP traps. Instead, you can disable or reenable the logging of OSPFv3-related events such as neighbor state changes and database overflow conditions. By default, the HP device logs these events.

To disable the logging of events, enter the following command:

```
ProCurveRS(config-ospf6-router)# no log-status-change
```

Syntax: [no] log-status-change

To reenable the logging of events, enter the following command:

```
ProCurveRS(config-ospf6-router)# log-status-change
```

Displaying OSPFv3 Information

You can display the information for the following OSPFv3 parameters:

- Areas
- Link state databases
- Interfaces
- Memory usage
- Neighbors
- Redistributed routes
- Routes
- SPF
- Virtual links
- Virtual neighbors

Displaying OSPFv3 Area Information

To display global OSPFv3 area information for the HP device, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf area
Area 0:
  Interface attached to this area: loopback 2 ethe 3/2 tunnel 2
  Number of Area scoped LSAs is 6
  Statistics of Area 0:
    SPF algorithm executed 16 times
    SPF last updated: 335256 sec ago
    Current SPF node count: 3
      Router: 2 Network: 1
      Maximum of Hop count to nodes: 2
...
```

Syntax: show ipv6 ospf area [<area-id>]

You can specify the <area-id> parameter in the following formats:

- As an IPv4 address, for example, 192.168.1.1.
- As a numerical value from 0 – 2,147,483,647.

The <area-id> parameter restricts the display to the specified OSPF area.

This display shows the following information:

Table 7.1: OSPFv3 area information fields

This Field...	Displays...
Area	The area number.
Interface attached to this area	The router interfaces attached to the area.
Number of Area scoped LSAs	Number of LSAs with a scope of the specified area.
SPF algorithm executed	The number of times the OSPF Shortest Path First (SPF) algorithm is executed within the area.
SPF last updated	The interval in seconds that the SPF algorithm was last executed within the area.
Current SPF node count	The current number of SPF nodes in the area.
Router	Number of router LSAs in the area.
Network	Number of network LSAs in the area.
Indx	The row number of the entry in the router's OSPF area table.
Area	The area number.
Maximum hop count to nodes.	The maximum number of hop counts to an SPF node within the area.

Displaying OSPFv3 Database Information

You can display a summary of the HP device's link state database or detailed information about a specified LSA type.

To display a summary of a device's link state database, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf database

Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Link 000001e6 223.223.223.223 800000ab 1547 8955   68
0               Link 000000d8 1.1.1.1         800000aa 1295 0639   68
0               Link 00000185 223.223.223.223 800000ab 1481 7e6b   56
0               Iap  00000077 223.223.223.223 800000aa 1404 966a   56
0               Rtr  00000124 223.223.223.223 800000b0 1397 912c   40
0               Net  00000016 223.223.223.223 800000aa 1388 1b09   32
0               Iap  000001d1 223.223.223.223 800000bd 1379 a072   72
0               Iap  000000c3 1.1.1.1         800000ae 1325 e021   52
0               Rtr  00000170 1.1.1.1         800000ad 1280 af8e   40
N/A             Extn 00000062 223.223.223.223 800000ae 1409 0ca7   32
N/A             Extn 0000021d 223.223.223.223 800000a8 1319 441e   32
```

Syntax: show ipv6 ospf database [advrtr <ipv4-address> | as-external | extensive | inter-prefix | inter-router | intra-prefix | link | link-id <number> | network | router [scope <area-id> | as | link]]

The **advrtr** <ipv4-address> parameter displays detailed information about the LSAs for a specified advertising router only.

The **as-external** keyword displays detailed information about the AS externals LSAs only.

The **extensive** keyword displays detailed information about all LSAs in the database.

The **inter-prefix** keyword displays detailed information about the inter-area prefix LSAs only.

The **inter-router** keyword displays detailed information about the inter-area router LSAs only.

The **intra-prefix** keyword displays detailed information about the intra-area prefix LSAs only.

The **link** keyword displays detailed information about the link LSAs only.

The **link-id** <number> parameter displays detailed information about the specified link LSAs only.

The **network** <number> displays detailed information about the network LSAs only.

The **router** <number> displays detailed information about the router LSAs only.

The **scope** <area-id> parameter displays detailed information about the LSAs for a specified area, AS, or link.

This display shows the following information:

Table 7.2: OSPFv3 database summary fields

This Field...	Displays...
Area ID	The OSPF area in which the HP device resides.
Type	Type of LSA. LSA types can be the following: • Rtr – Router LSAs (Type 1). • Net – Network LSAs (Type 2). • Inap – Inter-area prefix LSAs for ABRs (Type 3). • Inar – Inter-area router LSAs for ASBRs (Type 4). • Extn – AS external LSAs (Type 5). • Link – Link LSAs (Type 8). • Iap – Intra-area prefix LSAs (Type 9).
LS ID	The ID of the LSA, in hexadecimal, from which the device learned this route.
Adv Rtr	The device that advertised the route.
Seq(Hex)	The sequence number of the LSA. The OSPF neighbor that sent the LSA stamps it with a sequence number to enable the HP device and other OSPF routers to determine which LSA for a given route is the most recent.
Age	The age of the LSA, in seconds.
Chksum	A checksum for the LSA packet. The checksum is based on all the fields in the packet except the age field. The HP device uses the checksum to verify that the packet is not corrupted.
Len	The length, in bytes, of the LSA.

For example, to display detailed information about all LSAs in the database, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf database extensive
Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Link 00000031 1.1.1.1         80000001 35   6db9   56
     Router Priority: 1
     Options: V6E---R--
     LinkLocal Address: fe80::1
     Number of Prefix: 1
     Prefix Options:
     Prefix: 3002::/64
    ...
Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Iap  00000159 223.223.223.223 800000ab 357  946b   56
     Number of Prefix: 2
     Referenced LS Type: Network
     Referenced LS ID: 00000159
     Referenced Advertising Router: 223.223.223.223
     Prefix Options:  Metric: 0
     Prefix: 2000:4::/64
     Prefix Options:  Metric: 0
     Prefix: 2002:c0a8:46a::/64
Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Rtr  00000039 223.223.223.223 800000b1 355  8f2d   40
  Capability Bits: --E-
  Options: V6E---R--
  Type: Transit Metric: 1
  Interface ID: 00000058  Neighbor Interface ID: 00000058
  Neighbor Router ID: 223.223.223.223
Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Net  000001f4 223.223.223.223 800000ab 346  190a   32
     Options: V6E---R--
     Attached Router: 223.223.223.223
     Attached Router: 1.1.1.1
...
Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
N/A             Extn 000001df 223.223.223.223 800000af 368  0aa8   32
     Bits: E
     Metric: 00000001
     Prefix Options:
     Referenced LSType: 0
     Prefix: 2002::/16

Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
1               Inap 0000011d 10.1.1.188      80000001 124  25de   36
     Metric: 2
     Prefix Options:
     Prefix: 2000:2:2::/64

Area ID         Type LS ID    Adv Rtr         Seq(Hex) Age  Cksum  Len
0               Inar 0000005b 10.1.1.198      80000001 990  dbad   32
     Options: V6E---R--
     Metric: 1
     Destination Router ID:10.1.1.188
```

NOTE: Portions of this display are truncated for brevity. The purpose of this display is to show all possible fields that might display rather than to show complete output.

The fields that display depend upon the LSA type as shown in the following:

Table 7.3: OSPFv3 detailed database information fields

This Field...	Displays...
Router LSA (Type 1) (Rtr) Fields	
Capability Bits	A bit that indicates the capability of the HP device. The bit can be set to one of the following: • B – The device is an area border router. • E – The device is an AS boundary router. • V – The device is a virtual link endpoint. • W – The device is a wildcard multicast receiver.
Options	A 24-bit field that enables IPv6 OSPF routers to support the optional capabilities. When set, the following bits indicate the following: V6 – The device should be included in IPv6 routing calculations. E – The device floods AS-external-LSAs as described in RFC 2740. MC – The device forwards multicast packets as described in RFC 1586. N – The device handles type 7 LSAs as described in RFC 1584. R – The originator is an active router. DC –The device handles demand circuits.
Type	The type of interface. Possible types can be the following: • Point-to-point – A point-to-point connection to another router. • Transit – A connection to a transit network. • Virtual link – A connection to a virtual link.
Metric	The cost of using this router interface for outbound traffic.
Interface ID	The ID assigned to the router interface.
Neighbor Interface ID	The interface ID that the neighboring router has been advertising in hello packets sent on the attached link.
Neighbor Router ID	The router ID (IPv4 address) of the neighboring router that advertised the route. (By default, the HP router ID is the IPv4 address configured on the lowest numbered loopback interface. If the HP device does not have a loopback interface, the default router ID is the lowest numbered IPv4 address configured on the device.)

Table 7.3: OSPFv3 detailed database information fields (Continued)

This Field...	Displays...
Network LSA (Type 2) (Net) Fields	
Options	A 24-bit field that enables IPv6 OSPF routers to support the optional capabilities. When set, the following bits indicate the following: V6 – The device should be included in IPv6 routing calculations. E – The device floods AS-external-LSAs as described in RFC 2740. MC – The device forwards multicast packets as described in RFC 1586. N – The device handles type 7 LSAs as described in RFC 1584. R – The originator is an active router. DC –The device handles demand circuits.
Attached Router	The address of the neighboring router that advertised the route.
Inter-Area Prefix LSA (Type 3) (Inap) Fields	
Metric	The cost of the route.
Prefix Options	An 8-bit field describing various capabilities associated with the prefix.
Prefix	The IPv6 prefix included in the LSA.
Inter-Area Router LSA (Type 4) (Inar) Fields	
Options	A 24-bit field that enables IPv6 OSPF routers to support the optional capabilities. When set, the following bits indicate the following: V6 – The device should be included in IPv6 routing calculations. E – The device floods AS-external-LSAs as described in RFC 2740. MC – The device forwards multicast packets as described in RFC 1586. N – The device handles type 7 LSAs as described in RFC 1584. R – The originator is an active router. DC –The device handles demand circuits.
Metric	The cost of the route.
Destination Router ID	The ID of the router described in the LSA.
AS External LSA (Type 5) (Extn) Fields	
Bits	The bit can be set to one of the following: • E – If bit E is set, a Type 2 external metric. If bit E is zero, a Type 1 external metric. • F – A forwarding address is included in the LSA. • T – An external route tag is included in the LSA.
Metric	The cost of this route, which depends on bit E.
Prefix Options	An 8-bit field describing various capabilities associated with the prefix.

Table 7.3: OSPFv3 detailed database information fields (Continued)

This Field...	Displays...
Referenced LS Type	If non-zero, an LSA with this LS type is associated with the LSA.
Prefix	The IPv6 prefix included in the LSA.
Link LSA (Type 8) (Link) Fields	
Router Priority	The router priority of the interface attaching the originating router to the link.
Options	The set of options bits that the router would like set in the network LSA that will be originated for the link.
Link Local Address	The originating router's link-local interface address on the link.
Number of Prefix	The number of IPv6 address prefixes contained in the LSA.
Prefix Options	An 8-bit field of capabilities that serve as input to various routing calculations: • NU – The prefix is excluded from IPv6 unicast calculations. • LA – The prefix is an IPv6 interface address of the advertising router. • MC – The prefix is included in IPv6 multicast routing calculations. • P – NSSA area prefixes are readvertised at the NSSA area border.
Prefix	The IPv6 prefix included in the LSA.
Intra-Area Prefix LSAs (Type 9) (Iap) Fields	
Number of Prefix	The number of prefixes included in the LSA.
Referenced LS Type, Referenced LS ID	Identifies the router-LSA or network-LSA with which the IPv6 address prefixes are associated.
Referenced Advertising Router	The address of the neighboring router that advertised the route.
Prefix Options	An 8-bit field describing various capabilities associated with the prefix.
Metric	The cost of using the advertised prefix.
Prefix	The IPv6 prefix included in the LSA.
Number of Prefix	The number of prefixes included in the LSA.

Displaying OSPFv3 Interface Information

You can display a summary of information for all OSPFv3 interfaces or detailed information about a specified OSPFv3 interface.

To display a summary of OSPFv3 interfaces, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf interface

Interface  OSPF      Status State     Area
ethe 3/1             up
ethe 3/2   enabled   up     DR        0
ethe 3/4   disabled  down
loopback 2 enabled   up     Loopback  0
tunnel 1   disabled  down
tunnel 2   enabled   up     P2P       0
tunnel 6             up
```

Syntax: show ipv6 ospf interface [ethernet <port> | loopback <number> | tunnel <number> | ve <number>]

The **ethernet | loopback | tunnel | ve** parameter specifies the interface for which to display information. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a loopback, tunnel, or VE interface, also specify the number associated with the interface.

This display shows the following information:

Table 7.4: Summary of OSPFv3 interface information

This Field...	Displays...
Interface	The interface type, and the port number or number of the interface.
OSPF	The state of OSPFv3 on the interface. Possible states include the following: • Enabled. • Disabled.
Status	The status of the link. Possible status include the following: • Up. • Down.

Table 7.4: Summary of OSPFv3 interface information (Continued)

This Field...	Displays...
State	The state of the interface. Possible states includes the following: • DR – The interface is functioning as the Designated Router for OSPFv3. • BDR – The interface is functioning as the Backup Designated Router for OSPFv3. • Loopback – The interface is functioning as a loopback interface. • P2P – The interface is functioning as a point-to-point interface. • Passive – The interface is up but it does not take part in forming an adjacency. • Waiting – The interface is trying to determine the identity of the BDR for the network. • None – The interface does not take part in the OSPF interface state machine. • Down – The interface is unusable. No protocol traffic can be sent or received on such a interface. • DR other – The interface is a broadcast or NBMA network on which another router is selected to be the DR.
Area	The OSPF area to which the interface belongs.

For example, to display detailed information about Ethernet interface 2, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf interface ethernet 3/2
ethe 3/2 is up, type BROADCAST
  IPv6 Address:
      2002:c0a8:46a::1/64
      2000:4::106/64
  Instance ID 0, Router ID 223.223.223.223
  Area ID 0, Cost 1
  State DR, Transmit Delay 1 sec, Priority 1
  Timer intervals :
    Hello 10, Dead 40, Retransmit 5
  DR:223.223.223.223 BDR:1.1.1.1  Number of I/F scoped LSAs is 2
  DRElection:     5 times, DelayedLSAck:   523 times
  Neighbor Count = 1,   Adjacent Neighbor Count= 1
     Neighbor:
      1.1.1.1 (BDR)
    Statistics of interface ethe 3/2:
      Type       tx   rx tx-byte rx-byte
      Unknown     0    0       0       0
      Hello    3149 3138 1259284 1255352
      DbDesc      7    6     416     288
      LSReq       2    2      80     152
      LSUpdate 1508  530  109508   39036
      LSAck     526 1398   19036   54568
```

This display shows the following information:

Table 7.5: Detailed OSPFv3 interface information

This Field...	Displays...
Interface status	The status of the interface. Possible status includes the following: • Up. • Down.
Type	The type of OSPFv3 circuit running on the interface. Possible types include the following: • BROADCAST • POINT TO POINT UNKNOWN
IPv6 Address	The IPv6 address(es) assigned to the interface.
Instance ID	An identifier for an instance of OSPFv3.
Router ID	The IPv4 address of the HP device. By default, the HP router ID is the IPv4 address configured on the lowest numbered loopback interface. If the device does not have a loopback interface, the default router ID is the lowest numbered IPv4 address configured on the device.
Area ID	The IPv4 address or numerical value of the area in which the interface belongs.
Cost	The overhead required to send a packet through the interface.
State	The state of the interface. Possible states include the following: • DR – The interface is functioning as the Designated Router for OSPFv3. • BDR – The interface is functioning as the Backup Designated Router for OSPFv3. • Loopback – The interface is functioning as a loopback interface. • P2P – The interface is functioning as a point-to-point interface. • Passive – The interface is up but it does not take part in forming an adjacency. • Waiting – The interface is trying to determine the identity of the BDR for the network. • None – The interface does not take part in the OSPF interface state machine. • Down – The interface is unusable. No protocol traffic can be sent or received on such a interface. • DR other – The interface is a broadcast or NBMA network on which another router is selected to be the DR.
Transmit delay	The amount of time, in seconds, it takes to transmit Link State Updates packets on the interface.
Priority	The priority used when selecting the DR and the BDR. If the priority is 0, the interface does not participate in the DR and BDR election.

Table 7.5: Detailed OSPFv3 interface information (Continued)

This Field...	Displays...
Timer intervals	The interval, in seconds, of the hello-interval, dead-interval, and retransmit-interval timers.
DR	The router ID (IPv4 address) of the DR.
BDR	The router ID (IPv4 address) of the BDR.
Number of I/F scoped LSAs	The number of interface LSAs scoped for a specified area, AS, or link.
DR Election	The number of times the DR election occurred.
Delayed LSA Ack	The number of the times the interface sent a delayed LSA acknowledgement.
Neighbor Count	The number of neighbors to which the interface is connected.
Adjacent Neighbor Count	The number of neighbors with which the interface has formed an active adjacency.
Neighbor	The router ID (IPv4 address) of the neighbor. This field also identifies the neighbor as a DR or BDR, if appropriate.
Interface statistics	The following statistics are provided for the interface: • Unknown – The number of Unknown packets transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received Unknown packets. • Hello – The number of Hello packets transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received Hello packets. • DbDesc – The number of Database Description packets transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received Database Description packets. • LSReq – The number of link-state requests transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received link-state requests. • LSUpdate – The number of link-state updates transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received link-state requests. • LSAck – The number of link-state acknowledgements transmitted and received by the interface. Also, the total number of bytes associated with transmitted and received link-state acknowledgements.

Displaying OSPFv3 Memory Usage

To display information about OSPFv3 memory usage, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf memory
Total Static Memory Allocated : 5829 bytes
Total Dynamic Memory Allocated : 0 bytes
Memory Type               Size       Allocated  Max-alloc  Alloc-Fails
MTYPE_OSPF6_TOP           0          0          0          0
MTYPE_OSPF6_LSA_HDR       0          0          0          0
MTYPE_OSPF6_RMAP_COMPILED 0          0          0          0
MTYPE_OSPF6_OTHER         0          0          0          0
MTYPE_THREAD_MASTER       0          0          0          0
MTYPE_OSPF6_AREA          0          0          0          0
MTYPE_OSPF6_AREA_RANGE    0          0          0          0
MTYPE_OSPF6_SUMMARY_ADDRE 0          0          0          0
MTYPE_OSPF6_IF            0          0          0          0
MTYPE_OSPF6_NEIGHBOR      0          0          0          0
MTYPE_OSPF6_ROUTE_NODE    0          0          0          0
MTYPE_OSPF6_ROUTE_INFO    0          0          0          0
MTYPE_OSPF6_PREFIX        0          0          0          0
MTYPE_OSPF6_LSA           0          0          0          0
MTYPE_OSPF6_VERTEX        0          0          0          0
MTYPE_OSPF6_SPFTREE       0          0          0          0
MTYPE_OSPF6_NEXTHOP       0          0          0          0
MTYPE_OSPF6_EXTERNAL_INFO 0          0          0          0
MTYPE_THREAD              0          0          0          0
```

Syntax: show ipv6 ospf memory

This display shows the following information:

Table 7.6: OSPFv3 memory usage information

This Field...	Displays...
Total Static Memory Allocated	A summary of the amount of static memory allocated, in bytes, to OSPFv3.
Total Dynamic Memory Allocated	A summary of the amount of dynamic memory allocated, in bytes, to OSPFv3.
Memory Type	The type of memory used by OSPFv3. (This information is for use by HP's technical support in case of a problem.)
Size	The size of a memory type.
Allocated	The amount of memory currently allocated to a memory type.
Max-alloc	The maximum amount of memory that was allocated to a memory type.
Alloc-Fails	The number of times an attempt to allocate memory to a memory type failed.

Displaying OSPFv3 Neighbor Information

You can display a summary of OSPFv3 neighbor information for the HP device or detailed information about a specified neighbor.

To display a summary of OSPFv3 neighbor information for the device, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf neighbor
RouterID        Pri State    DR              BDR             Interface[State]
1.1.1.1           1 Full     223.223.223.223 1.1.1.1         ethe 3/2   [DR]
```

Syntax: show ipv6 ospf neighbor [router-id <ipv4-address>]

The **router-id** <ipv4-address> parameter displays only the neighbor entries for the specified router.

This display shows the following information:

Table 7.7: Summary of OSPFv3 neighbor information

Field	Description
Router ID	The IPv4 address of the neighbor. By default, the HP router ID is the IPv4 address configured on the lowest numbered loopback interface. If the device does not have a loopback interface, the default router ID is the lowest numbered IPv4 address configured on the device.
Pri	The OSPFv3 priority of the neighbor. The priority is used during election of the DR and BDR.
State	The state between the HP device and the neighbor. The state can be one of the following: • Down • Attempt • Init • 2-Way • ExStart • Exchange • Loading • Full
DR	The router ID (IPv4 address) of the DR.
BDR	The router ID (IPv4 address) of the BDR.

Table 7.7: Summary of OSPFv3 neighbor information (Continued)

Field	Description
Interface [State]	The interface through which the Routing Switch is connected to the neighbor. The state of the interface can be one of the following: • DR – The interface is functioning as the Designated Router for OSPFv3. • BDR – The interface is functioning as the Backup Designated Router for OSPFv3. • Loopback – The interface is functioning as a loopback interface. • P2P – The interface is functioning as a point-to-point interface. • Passive – The interface is up but it does not take part in forming an adjacency. • Waiting – The interface is trying to determine the identity of the BDR for the network. • None – The interface does not take part in the OSPF interface state machine. • Down – The interface is unusable. No protocol traffic can be sent or received on such a interface. • DR other – The interface is a broadcast or NBMA network on which another router is selected to be the DR.

For example, to display detailed information about a neighbor with the router ID of 1.1.1.1, enter the following command at any CLI level:

```
ProCurveRS# show ipv6 ospf neighbor router-id 3.3.3.3
RouterID        Pri State    DR              BDR             Interface[State]
3.3.3.3           1 Full     3.3.3.3         1.1.1.1         ve 10    [BDR]
DbDesc bit for this neighbor: --s
Nbr Ifindex of this router: 1
Nbr DRDecision: DR 3.3.3.3, BDR 1.1.1.1
Last received DbDesc: opt:xxx ifmtu:0 bit:--s seqnum:0
Number of LSAs in DbDesc retransmitting: 0
Number of LSAs in SummaryList: 0
Number of LSAs in RequestList: 0
Number of LSAs in RetransList: 0
SeqnumMismatch 0 times, BadLSReq 0 times
OnewayReceived 0 times, InactivityTimer 0 times
DbDescRetrans 0 times, LSReqRetrans 0 times
LSUpdateRetrans 1 times
LSAReceived 12 times, LSUpdateReceived 6 times
```

This display shows the following information:

Table 7.8: Detailed OSPFv3 neighbor information

Field	Description
Router ID	For information about this field, see Table 7.7 on page 7-29.
Pri	For information about this field, see Table 7.7 on page 7-29.
State	For information about this field, see Table 7.7 on page 7-29.

Table 7.8: Detailed OSPFv3 neighbor information (Continued)

Field	Description
DR	For information about this field, see Table 7.7 on page 7-29.
BDR	For information about this field, see Table 7.7 on page 7-29.
Interface [State]	For information about this field, see Table 7.7 on page 7-29.
DbDesc bit...	The Database Description packet, which includes 3 bits of information: • The first bit can be “i” or “-”. “i” indicates the inet bit is set. “-” indicates the inet bit is not set. • The second bit can be “m” or “-”. “m” indicates the more bit is set. “-” indicates the more bit is not set. • The third bit can be “m” or “s”. An “m” indicates the master. An “s” indicates standby.
Index	The ID of the LSA from which the neighbor learned of the router.
DR Decision	The router ID (IPv4 address) of the neighbor’s elected DR and BDR.
Last Received Db Desc	The content of the last database description received from the specified neighbor.
Number of LSAs in Db Desc retransmitting	The number of LSAs that need to be retransmitted to the specified neighbor.
Number of LSAs in Summary List	The number of LSAs in the neighbor’s summary list.
Number of LSAs in Request List	The number of LSAs in the neighbor’s request list.
Number of LSAs in Retransmit List	The number of LSAs in the neighbor’s retransmit list.
Seqnum Mismatch	The number of times sequence number mismatches occurred.
BadLSReq	The number of times the neighbor received a bad link-state request from the HP device.
One way received	The number of times a hello packet, which does not mention the router, is received from the neighbor. This omission in the hello packet indicates that the communication with the neighbor is not bidirectional.
Inactivity Timer	The number of times that the neighbor’s inactivity timer expired.
Db Desc Retransmission	The number of times sequence number mismatches occurred.
LSReqRetrans	The number of times the neighbor retransmitted link-state requests to the HP device.
LSUpdateRetrans	The number of times the neighbor retransmitted link-state updates to the HP device.
LSA Received	The number of times the neighbor received LSAs from the HP device.
LS Update Received	The number of times the neighbor received link-state updates from the HP device.

Displaying Routes Redistributed into OSPFv3

You can display all IPv6 routes or a specified IPv6 route that the HP device has redistributed into OSPFv3.

To display all IPv6 routes that the device has redistributed into OSPFv3, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf redistribute route
Id     Prefix                                    Protocol  Metric Type  Metric
 snIpAsPathAccessListStringRegExpression
1      2002::/16                                 Static    Type-2       1
2      2002:1234::/32                            Static    Type-2       1
```

Syntax: show ipv6 ospf redistribute route [<ipv6-prefix>]

The <ipv6-prefix> parameter specifies an IPv6 network prefix. (You do not need to specify the length of the prefix.)

For example, to display redistribution information for the prefix 2002::, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf redistribute route 2002::
Id     Prefix                                  Protocol  Metric Type  Metric
1      2002::/16                               Static    Type-2       1
```

These displays show the following information:

Table 7.9: OSPFv3 redistribution information

This Field...	Displays...
ID	An ID for the redistributed route.
Prefix	The IPv6 routes redistributed into OSPFv3.
Protocol	The protocol from which the route is redistributed into OSPFv3. Redistributed protocols can be the following: • BGP – BGP4+. • RIP – RIPng. • Static – IPv6 static route table. • Connected – A directly connected network.
Metric Type	The metric type used for routes redistributed into OSPFv3. The metric type can be the following: • Type-1 – Specifies a small metric (2 bytes). • Type-2 – Specifies a big metric (3 bytes).
Metric	The value of the default redistribution metric, which is the OSPF cost of redistributing the route into OSPFv3.

Displaying OSPFv3 Route Information

You can display the entire OSPFv3 route table for the HP device or only the route entries for a specified destination.

To display the entire OSPFv3 route table for the device, enter the following command at any level of the CLI:

```
  ProCurveRS# show ipv6 ospf routes
Current Route count: 4
    Intra: 4 Inter: 0 External: 0 (Type1 0/Type2 0)
    Equal-cost multi-path: 0
    Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 2000:4::/64                    V6E---R-- 0.0.0.0            1 0
    ::                             ethe 3/2
*IA 2002:c0a8:46a::/64             V6E---R-- 0.0.0.0            1 0
    ::                             ethe 3/2
*IA 2999::1/128                    --------- 0.0.0.0            0 0
    ::                             loopback 2
*IA 2999::2/128                    V6E---R-- 0.0.0.0            1 0
    fe80::2e0:52ff:fe91:bb37       ethe 3/2
```

Syntax: show ipv6 ospf routes [<ipv6-prefix>]

The <ipv6-prefix> parameter specifies a destination IPv6 prefix. (You do not need to specify the length of the prefix.) If you use this parameter, only the route entries for this destination are shown.

For example, to display route information for the destination prefix 2000:4::, enter the following command at any level of the CLI:

```
 ProCurveRS# show ipv6 ospf routes 2000:4::
 Destination                    Options   Area            Cost Type2 Cost
    Next Hop Router                Outgoing Interface
*IA 2000:4::/64                    V6E---R-- 0.0.0.0            1 0
    ::                             ethe 3/2
```

These displays show the following information:

Table 7.10: OSPFv3 route information

This Field...	Displays...
Current Route Count (Displays with the entire OSPFv3 route table only)	The number of route entries currently in the OSPFv3 route table.
Intra/Inter/External (Type1/Type2) (Displays with the entire OSPFv3 route table only)	The breakdown of the current route entries into the following route types: • Inter – The number of routes that pass into another area. • Intra – The number of routes that are within the local area. • External1 – The number of type 1 external routes. • External2 – The number of type 2 external routes.
Equal-cost multi-path (Displays with the entire OSPFv3 route table only)	The number of equal-cost routes to the same destination in the OSPFv3 route table. If load sharing is enabled, the Routing Switch equally distributes traffic among the routes.
Destination	The IPv6 prefixes of destination networks to which the HP device can forward IPv6 packets. “*IA” indicates the next Routing Switch is an intra-area router.

Table 7.10: OSPFv3 route information (Continued)

This Field...	Displays...
Options	A 24-bit field that enables IPv6 OSPF routers to support the optional capabilities. When set, the following bits indicate the following: V6 – The device should be included in IPv6 routing calculations. E – The device floods AS-external-LSAs as described in RFC 2740. MC – The device forwards multicast packets as described in RFC 1586. N – The device handles type 7 LSAs as described in RFC 1584. R – The originator is an active router. DC –The device handles demand circuits.
Area	The area whose link state information has led to the routing table entry's collection of paths.
Cost	The type 1 cost of this route.
Type2 Cost	The type 2 cost of this route.
Next-Hop Router	The IPv6 address of the next router a packet must traverse to reach a destination.
Outgoing Interface	The router interface through which a packet must traverse to reach the next-hop router.

Displaying OSPFv3 SPF Information

You can display the following OSPFv3 SPF information:

- SPF node information for a specified area.
- SPF table for a specified area.
- SPF tree for a specified area.

For example, to display information about SPF nodes in area 0, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf spf node area 0
SPF node for Area 0
SPF node 223.223.223.223,  cost: 0,  hops: 0
 nexthops to node:
 parent nodes:
 child nodes: 223.223.223.223:88

SPF node 223.223.223.223:88,  cost: 1,  hops: 1
 nexthops to node:    :: ethe 3/2
 parent nodes: 223.223.223.223
 child nodes: 1.1.1.1:0

SPF node 1.1.1.1:0,  cost: 1,  hops: 2
 nexthops to node:    fe80::2e0:52ff:fe91:bb37 ethe 3/2
 parent nodes: 223.223.223.223:88
 child nodes:
```

Syntax: show ipv6 ospf spf node area [<area-id>]

The **node** keyword displays SPF node information.

The **area** <area-id> parameter specifies a particular area. You can specify the <area-id> in the following formats:

- As an IPv4 address; for example, 192.168.1.1.
- As a numerical value from 0 – 2,147,483,647.

This display shows the following information:

Table 7.11: OSPFv3 SPF node information

This Field...	Displays...
SPF node	Each SPF node is identified by its router ID (IPv4 address). If the node is a child node, it is additionally identified by an interface on which the node can be reached appended to the router ID in the format <router-id>:<interface-id>.
Cost	The cost of traversing the SPF node to reach the destination.
Hops	The number of hops needed to reach the parent SPF node.
Next Hops to Node	The IPv6 address of the next hop-router and/or the router interface through which to access the next-hop router.
Parent Nodes	The SPF node's parent nodes. A ***parent node*** is an SPF node at the highest level of the SPF tree, which is identified by its router ID.
Child Nodes	The SPF node's child nodes. A ***child node*** is an SPF node at a lower level of the SPF tree, which is identified by its router ID and interface on which the node can be reached.

For example, to display the SPF table for area 0, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf spf table area 0
  SPF table for Area 0
  Destination         Bits Options  Cost  Nexthop                    Interface
R 1.1.1.1             ---- V6E---R-    1  fe80::2e0:52ff:fe91:bb37  ethe 3/2
N 223.223.223.223[88] ---- V6E---R-    1  ::                         ethe 3/2
```

Syntax: show ipv6 ospf spf table area <area-id>

The **table** parameter displays the SPF table.

The **area** <area-id> parameter specifies a particular area. You can specify the <area-id> in the following formats:

- As an IPv4 address, for example, 192.168.1.1.
- As a numerical value from 0 – 2,147,483,647.

This display shows the following information:

Table 7.12: OSPFv3 SPF Table

This Field...	Displays...
Destination	The destination of a route, which is identified by the following: • "R", which indicates the destination is a router. "N", which indicates the destination is a network. • An SPF node's router ID (IPv4 address). If the node is a child node, it is additionally identified by an interface on which the node can be reached appended to the router ID in the format <router-id>:<interface-id>.
Bits	A bit that indicates the capability of the HP device . The bit can be set to one of the following: • B – The device is an area border router. • E – The device is an AS boundary router. • V – The device is a virtual link endpoint. • W – The device is a wildcard multicast receiver.
Options	A 24-bit field that enables IPv6 OSPF routers to support the optional capabilities. When set, the following bits indicate the following: V6 – The router should be included in IPv6 routing calculations. E – The router floods AS-external-LSAs as described in RFC 2740. MC – The router forwards multicast packets as described in RFC 1586. N – The router handles type 7 LSAs as described in RFC 1584. R – The originator is an active router. DC –The router handles demand circuits.
Cost	The cost of traversing the SPF node to reach the destination.
Next hop	The IPv6 address of the next hop-router.

Table 7.12: OSPFv3 SPF Table (Continued)

This Field...	Displays...
Interface	The router interface through which to access the next-hop router.

For example, to display the SPF tree for area 0, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf spf tree area 0
   SPF tree for Area 0
 +- 223.223.223.223 cost 0
     +- 223.223.223.223:88 cost 1
         +- 1.1.1.1:0 cost 1
```

Syntax: show ipv6 ospf spf tree area <area-id>

The **tree** keyword displays the SPF table.

The **area** <area-id> parameter specifies a particular area. You can specify the <area-id> in the following formats:

- As an IPv4 address; for example, 192.168.1.1.
- As a numerical value from 0 – 2,147,483,647.

In this sample output, consider the SPF node with the router ID 223.223.223.223 to be the top (root) of the tree and the local router. Consider all other layers of the tree (223.223.223.223:88 and 1.1.1.1:0) to be destinations in the network. Therefore, traffic destined from router 223.223.223.223 to router 1.1.1.1:0 must first traverse router 223.223.223.223:88.

Displaying IPv6 OSPF Virtual Link Information

To display OSPFv3 virtual link information for the HP device, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf virtual-link
Index Transit Area ID  Router ID        Interface Address             State
1     1                1.1.1.1          3003::2                       P2P
```

Syntax: show ipv6 ospf virtual-link

This display shows the following information:

Table 7.13: OSPFv3 virtual link information

This Field...	Displays...
Index	An index number associated with the virtual link.
Transit Area ID	The ID of the shared area of two ABRs that serves as a connection point between the two routers.
Router ID	IPv4 address of the router at the other end of the virtual link (virtual neighbor).
Interface Address	The local address used to communicate with the virtual neighbor.

Table 7.13: OSPFv3 virtual link information (Continued)

This Field...	Displays...
State	The state of the virtual link. Possible states include the following: • P2P – The link is functioning as a point-to-point interface. • DOWN – The link is down.

Displaying OSPFv3 Virtual Neighbor Information

To display OSPFv3 virtual neighbor information for the HP device, enter the following command at any level of the CLI:

```
ProCurveRS# show ipv6 ospf virtual-neighbor
Index Router ID        Address                       State     Interface
1     1.1.1.1          3002::1                       Full      ethe 2/3
```

Syntax: show ipv6 ospf virtual-neighbor

This display shows the following information:

Table 7.14: OSPFv3 virtual neighbor information

This Field...	Displays...
Index	An index number associated with the virtual neighbor.
Router ID	IPv4 address of the virtual neighbor.
Address	The IPv6 address to be used for communication with the virtual neighbor.
State	The state between the HP device and the virtual neighbor. The state can be one of the following: • Down • Attempt • Init • 2-Way • ExStart • Exchange • Loading • Full
Interface	The IPv6 address of the virtual neighbor.