

IP Routing Features

Contents

Overview of IP Routing	5-6
IP Interfaces	5-7
IP Tables and Caches	5-7
ARP Cache Table	5-8
IP Route Table	5-8
IP Forwarding Cache	5-9
IP Route Exchange Protocols	5-10
IP Global Parameters for Routing Switches	5-10
ARP Age Timer	5-12
IP Interface Parameters for Routing Switches	5-14
Configuring IP Parameters for Routing Switches	5-15
Configuring IP Addresses	5-15
Changing the Router ID	5-15
Configuring ARP Parameters	5-17
How ARP Works	5-17
Enabling Proxy ARP	5-19
Enabling Local Proxy ARP	5-19
CLI Commands	5-20
Configuring Forwarding Parameters	5-21
Changing the TTL Threshold	5-21
Enabling Forwarding of Directed Broadcasts	5-21
Configuring ICMP	5-23
Disabling ICMP Messages	5-23
Disabling Replies to Broadcast Ping Requests	5-23
Disabling ICMP Destination Unreachable Messages	5-24
Disabling ICMP Redirects	5-25
Configuring Static IP Routes	5-25

Static Route Types	5-25
Other Sources of Routes in the Routing Table	5-26
Static IP Route Parameters	5-26
Static Route States Follow VLAN States	5-27
Configuring a Static IP Route	5-27
Displaying Static Route Information	5-29
Configuring the Default Route	5-29
Configuring RIP	5-30
Overview of RIP	5-30
RIP Parameters and Defaults	5-31
RIP Global Parameters	5-31
RIP Interface Parameters	5-31
Configuring RIP Parameters	5-32
Enabling RIP	5-32
Enabling IP RIP on a VLAN	5-33
Changing the RIP Type on a VLAN Interface	5-33
Changing the Cost of Routes Learned on a VLAN Interface	5-33
Configuring RIP Redistribution	5-34
Define RIP Redistribution Filters	5-34
Modify Default Metric for Redistribution	5-35
Enable RIP Route Redistribution	5-35
Changing the Route Loop Prevention Method	5-37
Displaying RIP Information	5-37
Displaying General RIP Information	5-38
Displaying RIP Interface Information	5-40
Displaying RIP Peer Information	5-41
Displaying RIP Redistribution Information	5-43
Displaying RIP Redistribution Filter (restrict) Information	5-43
Configuring OSPF	5-44
Terminology	5-45
Overview of OSPF	5-47
OSPF Router Types	5-48
Interior Routers	5-48
Area Border Routers (ABRs)	5-48
Autonomous System Boundary Router (ASBR)	5-49

Designated Routers	5-49
OSPF Area Types	5-51
Backbone Area	5-52
Normal Area	5-52
Not-So-Stubby-Area (NSSA)	5-53
Stub Area	5-54
OSPF RFC Compliance	5-54
Reducing AS External LSAs and Type-3 Summary LSAs	5-54
Algorithm for AS External LSA Reduction	5-55
Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA	5-56
Equal Cost Multi-Path Routing	5-57
Dynamic OSPF Activation and Configuration	5-59
General Configuration Steps for OSPF	5-60
Configuration Rules	5-61
OSPF Global and Interface Settings	5-61
Configuring OSPF on the Routing Switch	5-63
1. Enable IP Routing	5-63
2. Enable Global OSPF Routing	5-63
3. Changing the RFC 1583 OSPF Compliance Setting	5-64
4. Assign the Routing Switch to OSPF Areas	5-66
5. Assign VLANs and/or Subnets to Each Area	5-71
6. Optional: Assigning Loopback Addresses to an Area	5-72
7. Optional: Configure for External Route Redistribution in an OSPF Domain	5-74
8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone	5-77
9. Optional: Influence Route Choices by Changing the Administrative Distance Default	5-80
10. Optional: Change OSPF Trap Generation Choices	5-81
11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings	5-82
12. Optional: Configuring OSPF Interface Authentication	5-87
13. Configuring an ABR To Use a Virtual Link to the Backbone	5-89
Configuring a Virtual Link	5-90
Optional: Adjust Virtual Link Performance by Changing the Interface Settings	5-92

Configuring OSPF Authentication on a Virtual Link	5-95
OSPF Passive	5-97
Displaying OSPF Information	5-98
Displaying General OSPF Configuration Information	5-99
Displaying OSPF Area Information	5-100
Displaying OSPF External Link State Information	5-101
Displaying OSPF Interface Information	5-103
Displaying OSPF Interface Information for a Specific VLAN or IP Address	5-105
Displaying OSPF Link State Information	5-106
Displaying OSPF Neighbor Information	5-109
Displaying OSPF Redistribution Information	5-111
Displaying OSPF Redistribution Filter (restrict) Information	5-111
Displaying OSPF Virtual Neighbor Information	5-112
Displaying OSPF Virtual Link Information	5-113
Displaying OSPF Route Information	5-115
Displaying OSPF Traps Enabled	5-117
OSPF Equal-Cost Multipath (ECMP) for Different Subnets Available Through the Same Next-Hop Routes	5-117
Displaying the Current IP Load-Sharing Configuration	5-118
Configuring IRDP	5-120
Enabling IRDP Globally	5-121
Enabling IRDP on an Individual VLAN Interface	5-121
Displaying IRDP Information	5-122
Configuring DHCP Relay	5-123
Overview	5-123
DHCP Packet Forwarding	5-123
Unicast Forwarding	5-123
Broadcast Forwarding	5-124
Prerequisites for DHCP Relay Operation	5-124
Enabling DHCP Relay	5-124
Configuring an IP Helper Address	5-125
Operating Notes	5-125
Hop Count in DHCP Requests	5-125
Disabling the Hop Count in DHCP Requests	5-125

Operating Notes	5-126
Verifying the DHCP Relay Configuration	5-126
Displaying the DHCP Relay Setting	5-126
Displaying DHCP Helper Addresses	5-127
Displaying the Hop Count Setting	5-128
DHCP Option 82	5-128
Option 82 Server Support	5-129
Terminology	5-130
General DHCP Option 82 Requirements and Operation	5-131
Option 82 Field Content	5-132
Forwarding Policies	5-135
Configuration Options for Managing DHCP Client Request Packets	5-135
Multiple Option 82 Relay Agents in a Client Request Path	5-136
Validation of Server Response Packets	5-137
Multinetted VLANs	5-139
Configuring Option 82	5-139
Example of Option 82 Configuration	5-141
Operating Notes	5-142
UDP Broadcast Forwarding	5-144
Overview	5-144
Subnet Masking for UDP Forwarding Addresses	5-145
Configuring and Enabling UDP Broadcast Forwarding	5-146
Globally Enabling UDP Broadcast Forwarding	5-146
Configuring UDP Broadcast Forwarding on Individual VLANs	5-146
Displaying the Current IP Forward-Protocol Configuration	5-148
Operating Notes for UDP Broadcast Forwarding	5-149
Messages Related to UDP Broadcast Forwarding	5-149

Overview of IP Routing

The switches covered in this guide offer the following IP routing features, as noted:

- **IP Static Routes** – up to 256 static routes
- **RIP** (Router Information Protocol) – supports RIP Version 1, Version 1 compatible with Version 2 (default), and Version 2
- **OSPF** (Open Shortest Path First) – the standard routing protocol for handling larger routed networks
- **IRDP** (ICMP Router Discovery Protocol) – advertises the IP addresses of the routing interfaces on this switch to directly attached host systems
- **DHCP Relay** – allows you to extend the service range of your DHCP server beyond its single local network segment

License Requirements

In the 3500yl and 5400zl switches, OSPF is included with the Premium License. In the 6200yl and 8200zl switches, this feature is included with the base feature set.

Throughout this chapter, the switches covered in this guide are referred to as “routing switches”. When IP routing is enabled on your switch, it behaves just like any other IP router.

Basic IP routing configuration consists of adding IP addresses, enabling IP routing, and, enabling a route exchange protocol, such as Routing Information Protocol (RIP).

For configuring the IP addresses, refer to the chapter titled “Configuring IP Addresses” in the *Management and Configuration Guide* for your switch. The rest of this chapter describes IP routing and how to configure it in more detail. Use the information in this chapter if you need to change some of the IP parameters from their default values or you want to view configuration information or statistics.

IP Interfaces

On the routing switches, IP addresses are associated with individual VLANs. By default, there is a single VLAN (Default_VLAN) on the routing switch. In that configuration, a single IP address serves as the management access address for the entire device. If routing is enabled on the routing switch, the IP address on the single VLAN also acts as the routing interface.

Each IP address on a routing switch must be in a different subnet. You can have only one VLAN interface that is in a given subnet. For example, you can configure IP addresses 192.168.1.1/24 and 192.168.2.1/24 on the same routing switch, but you cannot configure 192.168.1.1/24 and 192.168.1.2/24 on the same routing switch.

You can configure multiple IP addresses on the same VLAN.

The number of IP addresses you can configure on an individual VLAN interface is 32.

You can use any of the IP addresses you configure on the routing switch for Telnet, Web management, or SNMP access, as well as for routing.

Note

All ProCurve devices support configuration and display of IP address in classical subnet format (example: 192.168.1.1 255.255.255.0) and Classless Interdomain Routing (CIDR) format (example: 192.168.1.1/24). You can use either format when configuring IP address information. IP addresses are displayed in classical subnet format only.

IP Tables and Caches

The following sections describe the IP tables and caches:

- ARP cache table
- IP route table
- IP forwarding cache

The software enables you to display these tables.

ARP Cache Table

The ARP cache contains entries that map IP addresses to MAC addresses. Generally, the entries are for devices that are directly attached to the routing switch.

An exception is an ARP entry for an interface-based static IP route that goes to a destination that is one or more router hops away. For this type of entry, the MAC address is either the destination device's MAC address or the MAC address of the router interface that answered an ARP request on behalf of the device, using proxy ARP.

ARP Cache. The ARP cache contains dynamic (learned) entries. The software places a dynamic entry in the ARP cache when the routing switch learns a device's MAC address from an ARP request or ARP reply from the device.

The software can learn an entry when the switch or routing switch receives an ARP request from another IP forwarding device or an ARP reply. Here is an example of a dynamic entry:

	IP Address	MAC Address	Type	Port
1	207.95.6.102	0800.5afc.ea21	Dynamic	6

Each entry contains the destination device's IP address and MAC address.

To configure other ARP parameters, see "Configuring ARP Parameters" on page 5-17.

IP Route Table

The IP route table contains routing paths to IP destinations.

Note

The default gateway, which you specify when you configure the basic IP information on the switch, is used only when routing is not enabled on the switch.

Routing Paths. The IP route table can receive the routing paths from the following sources:

- A directly-connected destination, which means there are no router hops to the destination
- A static IP route, which is a user-configured route
- A route learned through RIP
- A route learned through OSPF

Administrative Distance. The IP route table contains the best path to a destination. When the software receives paths from more than one of the sources listed above, the software compares the administrative distance of each path and selects the path with the lowest administrative distance. The administrative distance is a protocol-independent value from 1 – 255.

The IP route table is displayed by entering the CLI command **show ip route** from any context level in the console CLI. Here is an example of an entry in the IP route table:

Destination	Gateway	VLAN	Type	Sub-Type	Metric	D
10.10.10.1/32	10.10.12.1		connected		1	

Each IP route table entry contains the destination's IP address and subnet mask and the IP address of the next-hop router interface to the destination. Each entry also indicates route type, and for OSPF routes, the sub type, and the route's IP metric (cost). The type indicates how the IP route table received the route.

To configure a static IP route, see “Configuring a Static IP Route” on page 5-27

IP Forwarding Cache

The IP forwarding cache provides a fast-path mechanism for forwarding IP packets. The cache contains entries for IP destinations. When an ProCurve routing switch has completed processing and addressing for a packet and is ready to forward the packet, the device checks the IP forwarding cache for an entry to the packet's destination.

- If the cache contains an entry with the destination IP address, the device uses the information in the entry to forward the packet out the ports listed in the entry. The destination IP address is the address of the packet's final destination. The port numbers are the ports through which the destination can be reached.
- If the cache does not contain an entry, the software can create an entry in the forwarding cache.

Each entry in the IP forwarding cache has an age timer. The age interval depends on the number of entries in the table. The age timer ranges from 12 seconds (full table) to 36 seconds (empty table). Entries are only aged if they are not being utilized by traffic. If you have an entry that is always being used in hardware, it will never age. If there is no traffic, it will age in 12-36 seconds. The age timer is not configurable.

Note

You cannot add static entries to the IP forwarding cache.

IP Route Exchange Protocols

The switch supports the following IP route exchange protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

These protocols provide routes to the IP route table. You can use one or more of these protocols, in any combination. The protocols are disabled by default. For configuration information, see the following:

- “Configuring RIP” on page 5-30
- “Configuring OSPF” on page 5-44

IP Global Parameters for Routing Switches

The following table lists the IP global parameters and the page where you can find more information about each parameter.

Table 5-1. IP Global Parameters for Routing Switches

Parameter	Description	Default	See page
Router ID	The value that routers use to identify themselves to other routers when exchanging route information. OSPF uses the router ID to identify routers. RIP does not use the router ID.	The lowest-numbered IP address configured on the lowest-numbered routing interface.	5-15
Address Resolution Protocol (ARP)	A standard IP mechanism that routers use to learn the Media Access Control (MAC) address of a device on the network. The router sends the IP address of a device in the ARP request and receives the device’s MAC address in an ARP reply.	Enabled	5-17
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device’s ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age. (Can be set using the menu interface to be as long as 1440 minutes. Go to Menu > Switch Configuration > IP Config.) See “ARP Age Timer” on page 5-12.	Five minutes.	n/a

Parameter	Description	Default	See page
Proxy ARP	An IP mechanism a router can use to answer an ARP request on behalf of a host, by replying with the router's own MAC address instead of the host's.	Disabled	5-19
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	Refer to the chapter titled "Configuring IP Addressing" in the <i>Management and Configuration Guide</i> .
Directed broadcast forwarding	A directed broadcast is a packet containing all ones (or in some cases, all zeros) in the host portion of the destination IP address. When a router forwards such a broadcast, it sends a copy of the packet out each of its enabled IP interfaces. Note: You also can enable or disable this parameter on an individual interface basis. See table 5-2 on page 5-14.	Disabled	5-21
ICMP Router Discovery Protocol (IRDP)	An IP protocol that a router can use to advertise the IP addresses of its router interfaces to directly attached hosts. You can enable or disable the protocol at the Global CLI Config level. You also can enable or disable IRDP and configure the following protocol parameters on an individual VLAN interface basis at the VLAN Interface CLI Config level. <ul style="list-style-type: none"> • Forwarding method (broadcast or multicast) • Hold time • Maximum advertisement interval • Minimum advertisement interval • Router preference level 	Disabled	5-120 5-121
Static route	An IP route you place in the IP route table.	No entries	5-25
Default network route	The router uses the default network route if the IP route table does not contain a route to the destination. Enter an explicit default route (0.0.0.0 0.0.0.0 or 0.0.0.0/0) as a static route in the IP route table.	None configured	5-29

ARP Age Timer

The ARP age is the amount of time the switch keeps a MAC address learned through ARP in the ARP cache. The switch resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age.

You can increase the ARP age timeout maximum to 24 hours or more with this command:

Syntax: [no] ip arp-age <[1...1440] | infinite>

*Allows the ARP age to be set from 1 to 1440 minutes (24 hours). If the option “infinite” is configured, the internal ARP age timeout is set to 99,999,999 seconds (approximately 3.2 years). An **arp-age** value of 0 (zero) is stored in the configuration file to indicate that “infinite” has been configured. This value also displays with the **show** commands and in the menu display (**Menu > Switch Configuration > IP Config**).*

Default: 20 minutes.

```
ProCurve(config)# ip arp-age 1000
```

Figure 5-1. Example of Setting the ARP Age Timeout to 1000 Minutes

To view the value of ARP Age timer, enter the **show ip** command as shown in Figure 5-2.

```
ProCurve(config)# show ip

Internet (IP) Service

IP Routing : Disabled

Default Gateway : 15.255.120.1
Default TTL     : 64
Arp Age        : 1000
Domain Suffix   :
DNS server      :

VLAN            | IP Config  IP Address      Subnet Mask      Proxy ARP
-----+-----
DEFAULT_VLAN    | Manual     15.255.111.13   255.255.248.0    No
```

Figure 5-2. Example of show ip Command Displaying ARP Age

You can also view the value of the ARP Age timer in the configuration file.

```
ProCurve(config)# show running-config

Running configuration:

; J9091A Configuration Editor; Created on release #K.12.XX

hostname "8200LP"
module 2 type J8702A
module 3 type J8702A
module 4 type J8702A
ip default-gateway 15.255.120.1
[ip_arp_age_1000_ _]
snmp-server community "public" Unrestricted
snmp-server host 16.180.1.240 "public"
vlan 1
    name "DEFAULT_VLAN"
    untagged B1-B24,C1-C24,D1-D24
    ip address 15.255.120.85 255.255.248.0
    exit
gvrp
spanning-tree
```

Figure 5-3. Example Showing ip arp-age Value in the Running Config File

You can set or display the **arp-age** value using the menu interface (**Menu > Switch Configuration > IP Config**).

```

ProCurve                                     12-June-2007  14:45:31
=====-- TELNET - MANAGER MODE =====
                Switch Configuration - Internet (IP) Service

IP Routing : Disabled

Default Gateway : 15.255.120.1
Default TTL     : 64
Arp Age        : 1000

IP Config [Manual] : Manual

IP Address  : 15.255.111.11
Subnet Mask : 255.255.248.0

Actions->   Cancel   Edit   Save   Help
  
```

Figure 5-4. Example of the Menu Interface Displaying the ARP Age Value

IP Interface Parameters for Routing Switches

Table 5-2 lists the interface-level IP parameters for routing switches.

Table 5-2. IP Interface Parameters – Routing Switches

Parameter	Description	Default	See page
IP address	A Layer 3 network interface address; separate IP addresses on individual VLAN interfaces.	None configured	*
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	5-32
ICMP Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings. See table 5-1 on page 5-10 for global IRDP information.	Disabled	5-121
IP helper address	The IP address of a UDP application server (such as a BootP or DHCP server) or a directed broadcast address. IP helper addresses allow the routing switch to forward requests for certain UDP applications from a client on one subnet to a server on another subnet.	None configured	5-125

* Refer to the chapter titled “Configuring IP Addressing” in the *Management and Configuration Guide* for your switch.

Configuring IP Parameters for Routing Switches

The following sections describe how to configure IP parameters. Some parameters can be configured globally while others can be configured on individual VLAN interfaces. Some parameters can be configured globally and overridden for individual VLAN interfaces.

Note

This section describes how to configure IP parameters for routing switches. For IP configuration information when routing is not enabled, refer to the chapter titled “Configuring IP Addressing” in the *Management and Configuration Guide* for your routing switch.

Configuring IP Addresses

You can configure IP addresses on the routing switch’s VLAN interfaces. Configuring IP addresses is described in detail in the chapter titled “Configuring IP Addressing” in the *Management and Configuration Guide* for your switch.

Changing the Router ID

In most configurations, a routing switch has multiple IP addresses, usually configured on different VLAN interfaces. As a result, a routing switch’s identity to other devices varies depending on the interface to which the other device is attached. Some routing protocols, including Open Shortest Path First (OSPF), identify a routing switch by just one of the IP addresses configured on the routing switch, regardless of the interfaces that connect the routing switches. This IP address is the router ID.

Note

Routing Information Protocol (RIP) does not use the router ID.

If no router ID is configured, then, by default, the router ID on a ProCurve routing switch is the first IP address that becomes physically active at reboot. This is usually the lowest numbered IP interface configured on the device. However, if no router ID is configured and one or more user-configured loopback interfaces are detected at reboot, then the lowest-numbered (user-configured) loopback interface becomes the router ID. If the lowest-numbered loopback interface has multiple IP addresses, then the lowest of these addressees will be selected as the router ID. Once a router ID is selected, it will not automatically change unless a higher-priority interface is configured

IP Routing Features

Configuring IP Parameters for Routing Switches

on the routing switch *and* OSPF is restarted with a reboot. (User-Configured loopback interfaces are always higher priority than other configured interfaces.) However, you prefer, you can explicitly set the router ID to any valid IP address, as long as the IP address is not in use on another device in the network.

Note

To display the router ID, enter the **show ip ospf** CLI command at any Manager EXEC CLI level.

```
ProCurve(ospf)# show ip ospf

OSPF Configuration Information

 OSPF_protocol : enabled
 Router ID     : 10.10.10.1
-----

Currently defined areas:

Area ID          Type      Stub      Stub      Stub
-----          -
backbone         normal  1          send      ospf metric
0.0.0.2          nssa   10         send      external type 2
0.0.0.3          stub   2          send      ospf metric
0.0.0.4          stub   10         send      ospf metric
```

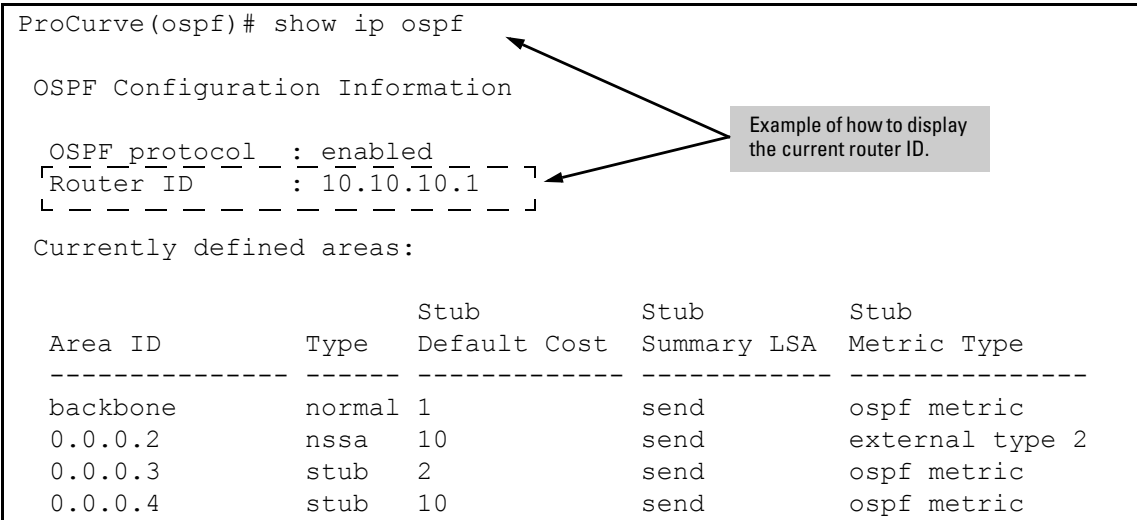


Figure 5-5. Example of show ip ospf Command with Router ID displayed

Reconfiguring the Router ID (Optional). If you want to change the router ID setting, do the following:

1. Go to the global config context. When you do so, the CLI prompt will appear similar to the following:
ProCurve(config)#_
2. If OSPF is not enabled, go to step 3. But if OSPF is enabled, then use **no router ospf** to disable OSPF operation.
3. Use **ip router-id <ip-addr>** to specify a new router ID. (This IP address must be unique in the routing switch configuration.)
4. If you disabled OSPF operation (step 2), then use **router ospf** to re-enable OSPF operation.

For more information on the router ID, refer to “IP Global Parameters for Routing Switches” on page 5-10 and “Changing the Router ID” on page 5-15.

To change the router ID, enter a command such as the following:

```
ProCurve(config)# ip router-id 209.157.22.26
```

Syntax: Syntax: ip router-id < ip-addr >

The < ip-addr > can be any valid, unique IP address.

Note

You can specify an IP address used for an interface on the ProCurve routing switch, but do not specify an IP address in use by another device.

Configuring ARP Parameters

Address Resolution Protocol (ARP) is a standard IP protocol that enables an IP routing switch to obtain the MAC address of another device’s interface when the routing switch knows the IP address of the interface. ARP is enabled by default and cannot be disabled.

How ARP Works

A routing switch needs to know a destination’s MAC address when forwarding traffic, because the routing switch encapsulates the IP packet in a Layer 2 packet (MAC layer packet) and sends the Layer 2 packet to a MAC interface on a device directly attached to the routing switch. The device can be the packet’s final destination or the next-hop router toward the destination.

The routing switch encapsulates IP packets in Layer 2 packets regardless of whether the ultimate destination is locally attached or is multiple router hops away. Since the routing switch’s IP route table and IP forwarding cache contain IP address information but not MAC address information, the routing switch cannot forward IP packets based solely on the information in the route table or forwarding cache. The routing switch needs to know the MAC address that corresponds with the IP address of either the packet’s locally attached destination or the next-hop router that leads to the destination.

For example, to forward a packet whose destination is multiple router hops away, the routing switch must send the packet to the next-hop router toward its destination, or to a default route or default network route if the IP route table does not contain a route to the packet’s destination. In each case, the

routing switch must encapsulate the packet and address it to the MAC address of a locally attached device, the next-hop router toward the IP packet's destination.

To obtain the MAC address required for forwarding a datagram, the routing switch does the following:

- First, the routing switch looks in the ARP cache (not the static ARP table) for an entry that lists the MAC address for the IP address. The ARP cache maps IP addresses to MAC addresses. The cache also lists the port attached to the device and, if the entry is dynamic, the age of the entry. A dynamic ARP entry enters the cache when the routing switch receives an ARP reply or receives an ARP request (which contains the sender's IP address and MAC address). A static entry enters the ARP cache from the static ARP table (which is a separate table) when the interface for the entry comes up.

To ensure the accuracy of the ARP cache, each dynamic entry has its own age timer. The timer is reset to zero each time the routing switch receives an ARP reply or ARP request containing the IP address and MAC address of the entry. If a dynamic entry reaches its maximum allowable age, the entry times out and the software removes the entry from the table. Static entries do not age out and can be removed only by you.

- If the ARP cache does not contain an entry for the destination IP address, the routing switch broadcasts an ARP request out all its IP interfaces. The ARP request contains the IP address of the destination. If the device with the IP address is directly attached to the routing switch, the device sends an ARP response containing its MAC address. The response is a unicast packet addressed directly to the routing switch. The routing switch places the information from the ARP response into the ARP cache.

ARP requests contain the IP address and MAC address of the sender, so all devices that receive the request learn the MAC address and IP address of the sender and can update their own ARP caches accordingly.

Note: The ARP request broadcast is a MAC broadcast, which means the broadcast goes only to devices that are directly attached to the routing switch. A MAC broadcast is not routed to other networks. However, some routers, including ProCurve routing switches, can be configured to reply to ARP requests from one network on behalf of devices on another network. See “Enabling Proxy ARP” below.

Note

If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP time-out and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

Enabling Proxy ARP

Proxy ARP allows a routing switch to answer ARP requests from devices on one network on behalf of devices in another network. Since ARP requests are MAC-layer broadcasts, they reach only the devices that are directly connected to the sender of the ARP request. Thus, ARP requests do not cross routers.

For example, if Proxy ARP is enabled on a routing switch connected to two subnets, 10.10.10.0/24 and 20.20.20.0/24, the routing switch can respond to an ARP request from 10.10.10.69 for the MAC address of the device with IP address 20.20.20.69. In standard ARP, a request from a device in the 10.10.10.0/24 subnet cannot reach a device in the 20.20.20.0 subnet if the subnets are on different network cables, and thus is not answered.

An ARP request from one subnet can reach another subnet when both subnets are on the same physical segment (Ethernet cable), since MAC-layer broadcasts reach all the devices on the segment.

Proxy ARP is disabled by default on ProCurve routing switches. To enable Proxy ARP, enter the following commands from the VLAN context level in the CLI:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip proxy-arp
```

To again disable IP proxy ARP, enter the following command:

```
ProCurve(vlan-1)# no ip proxy-arp
```

Syntax: [no] ip proxy-arp

Enabling Local Proxy ARP

When the Local Proxy ARP option is enabled, a switch responds with its MAC address to all ARP request on the VLAN. All IP packets are routed through and forwarded by the switch. The switch prevents broadcast ARP requests from reaching other ports on the VLAN.

Notes

Internet Control Message Protocol (ICMP) redirects will be disabled on interfaces on which local proxy ARP is enabled.

CLI Commands

To enable local proxy ARP, you must first enter vlan context, for example:

```
ProCurve(config) vlan 1
```

Then enter the command to enable local proxy ARP:

```
ProCurve(vlan-1) ip local-proxy-arp
```

Syntax: [no] ip local-proxy-arp

Enables the local proxy ARP option. You must be in VLAN context to execute this command. When enabled on a VLAN, the switch responds to all ARP requests received on the VLAN ports with its own hardware address.

*The **no** option disables the local proxy ARP option.*

Default: Disabled

Execute the **show ip** command to see which VLANs have local proxy ARP enabled.

```
ProCurve(vlan-1)# show ip

Internet (IP) Service

IP Routing : Disabled

Default TTL      : 64
Arp Age         : 20
Domain Suffix   :
DNS server      :

VLAN              | IP Config | IP Address      | Subnet Mask     | Proxy ARP
-----+-----+-----+-----+-----
DEFAULT_VLAN     | DHCP/Bootp | 15.255.157.54  | 255.255.248.0  | Yes Yes
VLAN2100         | Disabled
```

Figure 5-6. Local Proxy ARP is Enabled on the Default VLAN

Configuring Forwarding Parameters

The following configurable parameters control the forwarding behavior of ProCurve routing switches:

- Time-To-Live (TTL) threshold
- Forwarding of directed broadcasts

All these parameters are global and thus affect all IP interfaces configured on the routing switch.

To configure these parameters, use the procedures in the following sections.

Changing the TTL Threshold

The configuration of this parameter is covered in the chapter titled, “Configuring IP Addressing” in the *Management and Configuration Guide* for your routing switch.

Enabling Forwarding of Directed Broadcasts

A directed broadcast is an IP broadcast to all devices within a single directly-attached network or subnet. A net-directed broadcast goes to all devices on a given network. A subnet-directed broadcast goes to all devices within a given subnet.

Note

A less common type, the all-subnets broadcast, goes to all directly-attached subnets. Forwarding for this broadcast type also is supported, but most networks use IP multicasting instead of all-subnet broadcasting.

Forwarding for all types of IP directed broadcasts is disabled by default. You can enable forwarding for all types if needed. You cannot enable forwarding for specific broadcast types.

To enable forwarding of IP directed broadcasts, enter the following CLI command:

```
ProCurve (config) # ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

ProCurve software makes the forwarding decision based on the routing switch's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last-hop router.

IP Routing Features

Configuring IP Parameters for Routing Switches

To disable the directed broadcasts, enter the following CLI command:

```
ProCurve (config) # no ip directed-broadcast
```

Configuring ICMP

You can configure the following ICMP limits:

- **Burst-Normal** – The maximum number of ICMP replies to send per second.
- **Reply Limit** – You can enable or disable ICMP reply rate limiting.

Disabling ICMP Messages

ProCurve devices are enabled to reply to ICMP echo messages and send ICMP Destination Unreachable messages by default.

You can selectively disable the following types of Internet Control Message Protocol (ICMP) messages:

- **Echo messages** (ping messages) – The routing switch replies to IP pings from other IP devices.
- **Destination Unreachable messages** – If the routing switch receives an IP packet that it cannot deliver to its destination, the routing switch discards the packet and sends a message back to the device that sent the packet to the routing switch. The message informs the device that the destination cannot be reached by the routing switch.
- **Address Mask replies** – You can enable or disable ICMP address mask replies.

Disabling Replies to Broadcast Ping Requests

By default, ProCurve devices are enabled to respond to broadcast ICMP echo packets, which are ping requests. You can disable response to ping requests on a global basis using the following CLI method.

To disable response to broadcast ICMP echo packets (ping requests), enter the following command:

```
ProCurve(config)# no ip icmp echo broadcast-request
```

Syntax: [no] ip icmp echo broadcast-request

If you need to re-enable response to ping requests, enter the following command:

```
ProCurve(config)# ip icmp echo broadcast-request
```

Disabling ICMP Destination Unreachable Messages

By default, when a ProCurve device receives an IP packet that the device cannot deliver, the device sends an ICMP Unreachable message back to the host that sent the packet. The following types of ICMP Unreachable messages are generated:

- Administration – The packet was dropped by the ProCurve device due to a filter or ACL configured on the device.
- Fragmentation-needed – The packet has the “Don’t Fragment” bit set in the IP Flag field, but the ProCurve device cannot forward the packet without fragmenting it.
- Host – The destination network or subnet of the packet is directly connected to the ProCurve device, but the host specified in the destination IP address of the packet is not on the network.
- Network – The ProCurve device cannot reach the network specified in the destination IP address of the packet.
- Port – The destination host does not have the destination TCP or UDP port specified in the packet. In this case, the host sends the ICMP Port Unreachable message to the ProCurve device, which in turn sends the message to the host that sent the packet.
- Protocol – The TCP or UDP protocol on the destination host is not running. This message is different from the Port Unreachable message, which indicates that the protocol is running on the host but the requested protocol port is unavailable.
- Source-route-failure – The device received a source-routed packet but cannot locate the next-hop IP address indicated in the packet’s Source-Route option.

Note

Disabling an ICMP Unreachable message type does not change the ProCurve device’s ability to forward packets. Disabling ICMP Unreachable messages prevents the device from generating or forwarding the Unreachable messages.

To disable all ICMP Unreachable messages, enter the following command:

```
ProCurve(config)# no ip icmp unreachable
```

Syntax: [no] ip icmp unreachable

Disabling ICMP Redirects

You can disable ICMP redirects on the ProCurve routing switch only on a global basis, for all the routing switch interfaces. To disable ICMP redirects globally, enter the following command at the global CONFIG level of the CLI:

```
ProCurve(config)# no ip icmp redirects
```

Syntax: [no] ip icmp redirects

Configuring Static IP Routes

This feature enables you to create static routes (and null routes) by adding such routes directly to the route table. This section describes how to add static and null routes to the IP route table.

Static Route Types

You can configure the following types of static IP routes:

- **Standard** – the static route consists of a destination network address or host, a corresponding network mask, and the IP address of the next-hop IP address.
- **Null (discard)** – the Null route consists of the destination network address or host, a corresponding network mask, and either the **reject** or **blackhole** keyword. Typically, the null route is configured as a backup route for discarding traffic if the primary route is unavailable. By default, when IP routing is enabled, a route for the 127.0.0.0/8 network is created to the null interface. Traffic to this interface is rejected (dropped). This route is for all traffic to the “loopback” network, with the single exception of traffic to the host address of the switch’s loopback interface (127.0.0.1/32). Figure 5-8 on page 5-29 illustrates the default Null route entry in the switch’s routing table.

Note

On a single routing switch you can create one static route or null route to a given destination. Multiple static or null routes to the same destination are not supported.

Other Sources of Routes in the Routing Table

The IP route table can also receive routes from these other sources:

- **Directly-connected networks:** One route is created per IP interface. When you add an IP interface, the routing switch automatically creates a route for the network the interface is in.
- **RIP:** If RIP is enabled, the routing switch can learn about routes from the advertisements other RIP routers send to the routing switch. If the RIP route has a lower administrative distance than any other routes from different sources to the same destination, the routing switch places the route in the IP route table. (Refer to “Administrative Distance” on page 5-9.)
- **OSPF:** See RIP, but substitute “OSPF” for “RIP”.
- **Default route:** This is a specific static route that the routing switch uses if other routes to the destination are not available. See “Configuring the Default Route” on page 5-29.

Static IP Route Parameters

When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route’s destination network or host.
- The route’s path, which can be one of the following:
 - the IP address of a next-hop router.
 - a “null” interface. The routing switch drops traffic forwarded to the null interface.

The routing switch also applies default values for the following routing parameters:

- **The route’s metric:** In the case of static routes, this is the value the routing switch uses when comparing a static route to routes in the IP route table from other sources to the same destination. This is a fixed metric for static IP routes, and is set to “1”.
- **The route’s administrative distance (page 5-9):** In the case of static routes, this is the value the routing switch uses to compare a static route to routes from other route sources to the same destination before placing a route in the IP route table. The default administrative distance for static IP routes is 1, but can be configured to any value in the range of 1 - 255.

The fixed metric and administrative distance values ensure that the routing switch always prefers static IP routes over routes from other sources to the same destination.

Static Route States Follow VLAN States

IP static routes remain in the IP route table only so long as the IP interface to the next-hop router is up. If the next-hop interface goes down, the software removes the static route from the IP route table. If the next-hop interface comes up again, the software adds the route back to the route table.

This feature allows the routing switch to adjust to changes in network topology. The routing switch does not continue trying to use routes on unreachable paths but instead uses routes only when their paths are reachable.

For example, the following command configures a static route to 207.95.7.0 (with a network mask of 255.255.255.0), using 207.95.6.157 as the next-hop router's IP address.

```
ProCurve(config)# ip route 207.95.7.0/24 207.95.6.157
```

A static IP route specifies the route's destination address and the next-hop router's IP address or routing switch interface through which the routing switch can reach the destination. (The route is added to the routing switch's IP route table.)

In the above example, routing switch "A" knows that 207.95.6.157 is reachable through port A2, and assumes that local interfaces within that subnet are on the same port. Routing switch "A" deduces that IP interface 207.95.7.188 is also on port A2. The software automatically removes a static IP route from the route table if the next-hop VLAN used by that route becomes unavailable. When the VLAN becomes available again, the software automatically re-adds the route to the route table.

Configuring a Static IP Route

This feature includes these options:

- **Static Route:** configure a static route to a specific network or host address
- **Null Route:** configure a "null" route to discard IP traffic to a specific network or host address:
 - discard traffic for the destination, with ICMP notification to sender
 - discard traffic for the destination, without ICMP notification to sender

Syntax: [no] ip route < dest-ip-addr >/< mask-bits >
< next-hop-ip-addr | reject | blackhole > [distance]

dest-ip-addr >/< mask-bits: The route destination and network mask length for the destination IP address. Alternatively, you can enter the mask itself. For example, you can enter either **10.0.0.0/24** or **10.0.0.0 255.255.255.0** for a route destination of 10.0.0.0 255.255.255.0.

next-hop-ip-addr: This IP address is the gateway for reaching the destination. The next-hop IP address is not required to be directly reachable on a local subnet. (If the next-hop IP address is not directly reachable, the route will be added to the routing table as soon as a route to this address is learned.)

reject: Specifies a null route where IP traffic for the specified destination is discarded and an ICMP error notification is returned to the sender.

blackhole: Specifies a null route where IP traffic for the specified destination is discarded and no ICMP error notification is returned to the sender.

distance: Specifies the administrative distance to associate with a static route. If not specified, this value is set to a default of 1. For more on this topic, refer to “Administrative Distance” on page 5-9. (Range: 1 - 255)

The **no** form of the command deletes the specified route for the specified destination next-hop pair.

The following example configures two static routes for traffic delivery and identifies two other null routes for which traffic should be discarded instead of forwarded.

```
ProCurve(config)# ip route 10.10.40.0/24 10.10.10.1
ProCurve(config)# ip route 10.10.50.128/27 10.10.10.1
ProCurve(config)# ip route 10.10.20.177/32 reject
ProCurve(config)# ip route 10.10.30.0/24 blackhole
```

Configures static routes to two different network destinations using the same next-hop router IP address.

Configures a null route to drop traffic for the device at 10.50.10.177 and return an ICMP notification to the sender.

Configures a null route to drop traffic for the 10.50.10.0 network without any ICMP notification to the sender.

Figure 5-7. Examples of Configuring Static Routes

Displaying Static Route Information

The **show ip route static** command displays the current static route configuration on the routing switch. Figure 5-8 shows the configuration resulting from the static routes configured in the preceding example.

```
ProCurve(config)# show ip route static
```

IP Route Entries					
Destination	Gateway	VLAN	Type	Sub-Type	Metric Dist.
10.10.20.177/32	reject		static		1 1
10.10.40.0/24	VLAN10	10	static		1 1
10.10.50.128/27	VLAN10	10	static		1 1
10.11.30.0/24	blackhole		static		1 1
127.0.0.0/8	reject		static		0 0

This reject (default null) route is included by default.
Refer to "Static Route Types" on page 5-25

Figure 5-8. Example of Displaying the Currently Configured Static Routes

Configuring the Default Route

You can also assign the default route and enter it in the routing table. The default route is used for all traffic that has a destination network not reachable through any other IP routing table entry. For example, if 208.45.228.35 is the IP address of your ISP router, all non-local traffic could be directed to the ISP by entering this command:

```
ProCurve(config)# ip route 0.0.0.0/0 208.45.228.35
```

Configuring RIP

This section describes how to configure RIP using the CLI interface.

To display RIP configuration information and statistics, see “Displaying RIP Information” on page 5-37.

Overview of RIP

Routing Information Protocol (RIP) is an IP route exchange protocol that uses a *distance vector* (a number representing distance) to measure the cost of a given route. The cost is a distance vector because the cost often is equivalent to the number of router hops between the ProCurve routing switch and the destination network.

A ProCurve routing switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the IP route table as the route to the destination. Typically, the best path is the path with the fewest hops. A hop is another router through which packets must travel to reach the destination. If the ProCurve routing switch receives a RIP update from another router that contains a path with fewer hops than the path stored in the ProCurve routing switch's route table, the routing switch replaces the older route with the newer one. The routing switch then includes the new path in the updates it sends to other RIP routers, including ProCurve routing switches.

RIP routers, including ProCurve routing switches, also can modify a route's cost, generally by adding to it, to bias the selection of a route for a given destination. In this case, the actual number of router hops may be the same, but the route has an administratively higher cost and is thus less likely to be used than other, lower-cost routes. A RIP route can have a maximum cost of 15. Any destination with a higher cost is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

The switches covered in this guide support the following RIP types:

- Version 1
- V1 compatible with V2
- Version 2 (the default)

Note

ICMP Host Unreachable Message for Undeliverable ARPs. If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

RIP Parameters and Defaults

The following tables list the RIP parameters, their default values, and where to find configuration information.

RIP Global Parameters

5-3 lists the global RIP parameters and their default values.

Table 5-3. RIP Global Parameters

Parameter	Description	Default
RIP state	Routing Information Protocol V2-only.	Disabled
auto-summary	Enable/Disable advertisement of summarized routes.	Enabled
metric	Default metric for imported routes.	1
redistribution	RIP can redistribute static, connected, and OSPF routes. (RIP redistributes connected routes by default, when RIP is enabled.)	Disabled

RIP Interface Parameters

5-4 lists the VLAN interface RIP parameters and their default values.

Table 5-4. RIP Interface Parameters

Parameter	Description	Default
RIP version	The version of the protocol that is supported on the interface. The version can be one of the following: <ul style="list-style-type: none"> • Version 1 only • Version 2 only • Version 1 or version 2 	V2-only

Parameter	Description	Default
metric	A numeric cost the routing switch adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1
IP address	The routes that a routing switch learns or advertises can be controlled.	The routing switch learns and advertises all RIP routes on all RIP interfaces
loop prevention	The method the routing switch uses to prevent routing loops caused by advertising a route on the same interface as the one on which the routing switch learned the route. <ul style="list-style-type: none">• Split horizon - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.• Poison reverse - the routing switch assigns a cost of 16 ("infinite" or "unreachable") to a route before advertising it on the same interface as the one on which the routing switch learned the route.	Poison reverse
receive	Define the RIP version for incoming packets	V2-only
send	Define the RIP version for outgoing packets	V2-only

Configuring RIP Parameters

Use the following procedures to configure RIP parameters on a system-wide and individual VLAN interface basis.

Enabling RIP

RIP is disabled by default. To enable it, use one of the following methods. When you enable RIP, the default RIP version is **RIPv2-only**. You can change the RIP version on an individual interface basis to **RIPv1** or **RIPv1-or-v2** if needed.

To enable RIP on a routing switch, enter the following commands:

```
ProCurve(config)# ip routing
ProCurve(config)# router rip
ProCurve(rip)# exit
ProCurve(config)# write memory
```

Syntax: [no] router rip

Note

IP routing must be enabled prior to enabling RIP. The first command in the preceding sequence enables IP routing.

Enabling IP RIP on a VLAN

To enable RIP on all IP addresses in a VLAN, use **ip rip** in the VLAN context. When the command is entered without specifying any IP address, it is enabled in all configured IP addresses of the VLAN.

To enable RIP on a specific IP address in a VLAN, use **ip rip [*ip-addr*>|all]** in the VLAN context and enter a specific IP address. If you want RIP enabled on all IP addresses, you can specify **all** in the command instead of a specific IP address.

Changing the RIP Type on a VLAN Interface

When you enable RIP on a VLAN interface, **RIPv2-only** is enabled by default. You can change the RIP type to one of the following on an individual VLAN interface basis:

- Version 1 only
- Version 2 only (the default)
- Version 1 - or - version 2

To change the RIP type supported on a VLAN interface, enter commands such as the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip rip v1-only
ProCurve(vlan-1)# exit
ProCurve(config)# write memory
```

Syntax: [no] ip rip <v1-only | v1-or-v2 | v2-only >

Changing the Cost of Routes Learned on a VLAN Interface

By default, the switch interface increases the cost of a RIP route that is learned on the interface. The switch increases the cost by adding one to the route's metric before storing the route.

You can change the amount that an individual VLAN interface adds to the metric of RIP routes learned on the interface.

Note

RIP considers a route with a metric of 16 to be unreachable. Use this metric only if you do not want the route to be used. In fact, you can prevent the switch from using a specific interface for routes learned through that interface by setting its metric to 16.

To increase the cost a VLAN interface adds to RIP routes learned on that interface, enter commands such as the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip rip metric 5
```

These commands configure vlan-1 to add 5 to the cost of each route learned on the interface.

Syntax: ip rip metric < 1-16 >

Configuring RIP Redistribution

You can configure the routing switch to redistribute connected, static, and OSPF routes into RIP. When you redistribute a route into RIP, the routing switch can use RIP to advertise the route to its RIP neighbors.

To configure redistribution, perform the following tasks:

1. Configure redistribution filters to permit or deny redistribution for a route based on the destination network address or interface. (optional)
2. Enable redistribution

Define RIP Redistribution Filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On the switches covered in this guide, redistribution is supported for static routes, directly connected routes, and OSPF routes. Redistribution of any other routing protocol into RIP is not currently supported. When you configure redistribution for RIP, you can specify that static, connected, or OSPF routes are imported into RIP routes. Likewise, OSPF redistribution supports the import of static, connected, or RIP routes into OSPF routes.

To configure for redistribution, define the redistribution tables with “restrict” redistribution filters. In the CLI, use the **restrict** command for RIP at the RIP router level.

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Example: To configure the switch to filter out redistribution of static, connected, or OSPF routes on network 10.0.0.0, enter the following commands:

```
ProCurve(config)# router rip
ProCurve(rip)# restrict 10.0.0.0 255.0.0.0
ProCurve(rip)# write memory
```

Note

The default configuration permits redistribution for all default connected routes only.

Syntax: `restrict < ip-addr > < ip-mask > | < ip-addr /< prefix length >`

This command prevents any routes with a destination address that is included in the range specified by the address/mask pair from being redistributed by RIP.

Modify Default Metric for Redistribution

The default metric is a global parameter that specifies the cost applied to all RIP routes by default. The default value is 1. You can assign a cost from 1 – 15.

Example: To assign a default metric of 4 to all routes imported into RIP, enter the following commands:

```
ProCurve(config)# router rip
ProCurve(rip)# default-metric 4
```

Syntax: `default-metric < value >`

The `< value >` can be from 1 – 15. The default is 1.

Enable RIP Route Redistribution

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

To enable redistribution of connected and static IP routes into RIP, enter the following commands.

```
0 (config) # router rip
ProCurve (rip) # redistribute connected
ProCurve (rip) # redistribute static
ProCurve (rip) # redistribute ospf
ProCurve (rip) # write memory
```

Syntax: [no] redistribute < connected | static | ospf >

Changing the Route Loop Prevention Method

RIP can use the following methods to prevent routing loops:

- **Split horizon** - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.
- **Poison reverse** - the routing switch assigns a cost of 16 (“infinity” or “unreachable”) to a route before advertising it on the same interface as the one on which the routing switch learned the route. This is the default.

These loop prevention methods are configurable on an individual VLAN interface basis.

Note

These methods are in addition to RIP's maximum valid route cost of 15.

Poison reverse is enabled by default. Disabling poison reverse causes the routing switch to revert to **Split horizon**. (Poison reverse is an extension of Split horizon.) To disable Poison reverse on an interface, and thereby enable Split horizon, enter the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# no ip rip poison-reverse
```

Syntax: [no] ip rip poison-reverse

Entering the command without the “no” option will re-enable Poison reverse.

Displaying RIP Information

All RIP configuration and status information is shown by the CLI command **show ip rip** and options off that command. The following RIP information can be displayed:

RIP Information Type	Page
General Information	5-38
Interface Information	5-40
Peer Information	5-41
Redistribute Information	5-43
Restrict Information	5-43

Displaying General RIP Information

To display general RIP information, enter **show ip rip** at any context level. The resulting display will appear similar to the following:

```
ProCurve(config)# show ip rip

RIP global parameters

RIP protocol      : enabled
Auto-summary     : enabled
Default Metric   : 4
Distance         : 120
Route changes    : 0
Queries          : 0

RIP interface information

IP Address      Status      Send mode      Recv mode      Metric      Auth
-----
100.1.0.1      enabled    V2-only        V2-only        5           none
100.2.0.1      enabled    V2-only        V2-only        5           none
100.3.0.1      enabled    V2-only        V2-only        5           none
100.4.0.1      enabled    V2-only        V2-only        5           none
100.10.0.1     enabled    V2-only        V2-only        5           none
100.11.0.1     enabled    V2-only        V2-only        5           none
100.12.0.1     enabled    V2-only        V2-only        5           none

RIP peer information

IP Address      Bad routes  Last update timeticks
-----
```

Figure 5-9. Example of General RIP Information Listing

The display is a summary of Global RIP information, information about interfaces with RIP enabled, and information about RIP peers. The following fields are displayed:

- **RIP protocol** – Status of the RIP protocol on the router. RIP must be enabled here and on the VLAN interface for RIP to be active. The default is **disabled**.
- **Auto-summary** – Status of Auto-summary for all interfaces running RIP. If auto-summary is enabled, then subnets will be summarized to a class network when advertising outside of the given network.
- **Default Metric** – Sets the default metric for imported routes. This is the metric that will be advertised with the imported route to other RIP peers. A RIP metric is a measurement used to determine the 'best' path to network; 1 is the best, 15 is the worse, 16 is unreachable.
- **Route changes** – The number of times RIP has modified the routing switch's routing table.

- **Queries** – The number of RIP queries that have been received by the routing switch.
- **RIP Interface Information** – RIP information on the VLAN interfaces on which RIP is enabled.
 - **IP Address** – IP address of the VLAN interface running rip.
 - **Status** – Status of RIP on the VLAN interface.
 - **Send mode** – The format of the RIP updates: RIP 1, RIP 2, or RIP 2 version 1 compatible.
 - **Recv mode** – The switch can process RIP 1, RIP 2, or RIP 2 version 1 compatible update messages.
 - **Metric** – The path “cost”, a measurement used to determine the 'best' RIP route path; 1 is the best, 15 is the worse, 16 is unreachable.
 - **Auth** – RIP messages can be required to include an authentication key if enabled on the interface.
- **RIP Peer Information** – RIP Peers are neighboring routers from which the routing switch has received RIP updates.
 - **IP Address** – IP address of the RIP neighbor.
 - **Bad routes** – The number of route entries which were not processed for any reason.
 - **Last update timeticks** – How many seconds have passed since we received an update from this neighbor.

Syntax: show ip rip

Displaying RIP Interface Information

To display RIP interface information, enter the `show ip rip interface` command at any context level. The resulting display will appear similar to the following:

```
ProCurve# show ip rip interface
```

RIP interface information						
IP Address	Status	Send mode	Recv mode	Metric	Auth	
100.1.0.1	enabled	V2-only	V2-only	1	none	
100.2.0.1	enabled	V2-only	V2-only	1	none	
100.3.0.1	enabled	V2-only	V2-only	1	none	
100.4.0.1	enabled	V2-only	V2-only	1	none	

Figure 5-10. Example of Show IP RIP Interface Output

See “RIP Interface Information” on the previous page for definitions of these fields.

You can also display the information for a single RIP VLAN interface, by specifying the VLAN ID for the interface, or specifying the IP address for the interface.

Displaying RIP interface information by VLAN ID: For example, to show the RIP interface information for VLAN 1000, use the `show ip rip interface vlan <vid>` command.

```
ProCurve# show ip rip interface vlan 4
```

RIP configuration and statistics for VLAN 4	
RIP interface information for 100.4.0.1	
IP Address :	100.4.0.1
Status :	enabled
Send mode :	V2-only
Recv mode :	V2-only
Metric :	1
Auth :	none
Bad packets received :	0
Bad routes received :	0
Sent updates :	0

Figure 5-11. Example of RIP Interface Output by VLAN

The information in this display includes the following fields, which are defined under “RIP Interface Information” on page 5-39: **IP Address**, **Status**, **Send mode**, **Recv mode**, **Metric**, and **Auth**.

The information also includes the following fields:

- **Bad packets received** – The number of packets that were received on this interface and were not processed for any reason.
- **Bad routes received** – The number of route entries that were received on this interface and were not processed for any reason.
- **Sent updates** – The number of RIP routing updates that have been sent on this interface.

Displaying RIP interface information by IP Address: For example, to show the RIP interface information for the interface with IP address 100.2.0.1, enter the **show ip rip interface** command as shown below:

```
ProCurve# show ip rip interface 100.2.0.1
RIP interface information for 100.2.0.1
  IP Address : 100.2.0.1
  Status     : enabled
  Send mode  : V2-only
  Recv mode  : V2-only
  Metric     : 1
  Auth      : none
  Bad packets received : 0
  Bad routes received  : 0
  Sent updates        : 0
```

Figure 5-12. Example of Show IP RIP Interface Output by IP Address

The information shown in this display has the same fields as for the display for a specific VLAN ID. See the previous page for the definitions of these fields.

Syntax: show ip rip interface [*ip-addr* | vlan < *vlan-id* >]

Displaying RIP Peer Information

To display RIP peer information, enter the **show ip rip peer** command at any context level.

The resulting display will appear similar to the following:

```
ProCurve# show ip rip peer

RIP peer information

  IP Address          Bad routes  Last update timeticks
  -----
100.1.0.100          0           1
100.2.0.100          0           0
100.3.0.100          0           2
100.10.0.100         0           1
```

Figure 5-13. Example of Show IP RIP Peer Output

This display lists all neighboring routers from which the routing switch has received RIP updates. The following fields are displayed:

- **IP Address** – IP address of the RIP peer neighbor.
- **Bad routes** – The number of route entries that were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this peer neighbor.

Displaying RIP information for a specific peer: For example, to show the RIP peer information for the peer with IP address 100.1.0.100, enter **show ip rip peer 100.1.0.100**.

```
ProCurve# show ip rip peer 100.0.1.100

RIP peer information for 100.0.1.100

  IP Address : 100.1.0.100

  Bad routes : 0

  Last update timeticks : 2
```

Figure 5-14. Example of Show IP RIP Peer < ip-addr > Output

This display lists the following information for a specific RIP peer:

- **IP Address** – IP address of the RIP peer neighbor.
- **Bad routes** – The number of route entries which were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this neighbor.

Displaying RIP Redistribution Information

To display RIP redistribution information, enter the **show ip rip redistribute** command at any context level:

```
ProCurve# show ip rip redistribute

RIP redistributing

Route type Status
-----
connected enabled
static      disabled
ospf       disabled
```

Figure 5-15. Example of Show IP RIP Redistribute Output

RIP automatically redistributes connected routes that are configured on interfaces that are running RIP, and all routes that are learned via RIP. The **router rip redistribute** command, described on page 5-34, configures the routing switch to cause RIP to advertise connected routes that are not running RIP, static routes, and OSPF routes. The display shows whether RIP redistribution is enabled or disabled for connected, static, and OSPF routes.

Displaying RIP Redistribution Filter (restrict) Information

To display RIP restrict filter information, enter the **show ip rip restrict** command at any context level:

```
ProCurve# show ip rip restrict

RIP restrict list

IP Address      Mask
-----
```

Figure 5-16. Example of Show IP RIP Restrict Output

The display shows if any routes, identified by the IP Address and Mask fields are being restricted from redistribution. The restrict filters are configured by the **router ip restrict** command described on page 5-34.

Configuring OSPF

Feature	Default	Page
Enable IP Routing and Global OSPF Routing	disabled	5-63
Changing the RFC 1583 OSPF Compliance Setting	enabled	5-64
Assign the Routing Switch to OSPF Areas	n/a	5-66
Assign VLANs and/or Subnets to Each Area	n/a	5-71
External Route Redistribution	disabled	5-74
Configure Ranges on an ABR To Reduce Advertising	n/a	5-77
Use Administrative Distance To Influence Route Choices		5-80
Generate OSPF Traps	enabled	5-81
Cost Per Interface	1	5-83
Dead Interval Per Interface	40 sec.	5-83
Hello Interval Per Interface	10 sec.	5-84
Priority Per interface	1	5-84
Retransmit Interval Per Interface	5 sec.	5-85
Transit Delay Per Interface	1 sec.	5-85
Password and MD5 Authentication	disabled	5-87, 5-88
Virtual Link Configuration	n/a	5-89
Dead Interval on a Virtual Link	40 sec.	5-92
Hello Interval on a Virtual Link	10 sec.	5-93
Retransmit Interval on a Virtual Link	5 sec.	5-93
Transit Delay on a Virtual Link	1 sec.	5-94
Password and MD5 Authentication on a Virtual Link	disabled	5-95, 5-96
Displaying OSPF Information	n/a	5-98

This section describes how to configure OSPF using the CLI interface.

Terminology

Area Border Router (ABR): An OSPF-enabled router having interfaces on two or more OSPF areas. (Refer to “Area Border Routers (ABRs)” on page 5-48.)

Autonomous System (AS): A single interior gateway protocol (IGP) domain such as an OSPF or RIP domain.

Autonomous System Boundary Router (ASBR): An OSPF-enabled router having interfaces in multiple IGP domains, such as an ASBR with membership in both a normal area of an OSPF domain and a RIP domain. (Refer to “Autonomous System Boundary Router (ASBR)” on page 5-49.)

Backbone Area: Required in any OSPF domain, this is the transit area for all advertisements and routed traffic between non-backbone areas. (Refer to “Backbone Area” on page 5-52.)

Backup Designated Router (BDR): If the DR for a network becomes inaccessible, the BDR takes over the DR function. (See also “Designated Router”, below, and refer to “Designated Routers” on page 5-49.)

Default Route: A route defined as 0.0.0.0/0. OSPF uses type 3 (summary) defaults and type 7 (external) default routes.

Designated Router (DR): Used in networks having two or more routers, and serves as the distribution point for forwarding updates throughout the network. (See also “Backup Designated Router”, above, and refer to “Designated Routers” on page 5-49.)

External Type-5 Link-State Advertisement: An LSA summarizing known external links for the backbone and normal areas. Refer to Table 5-5 on page 5-47. (See also “Link State Advertisement”.)

External Type-7 Link State Advertisement: An LSA originating with an ASBR in an NSSA and allowed only in the NSSA. Refer to Table 5-5 on page 5-47. (See also “Link State Advertisement”.)

Interior Gateway Protocol (IGP): A method for forwarding traffic between autonomous routing domains. Commonly used between OSPF and RIP domains.

Interior Router: An OSPF-enabled routing switch having interfaces in only one OSPF area. (Refer to “Interior Routers” on page 5-48.)

Link-State Advertisement (LSA): A message sent by a router to its neighbors to advertise the existence of a route to a destination known by the originating router. Refer to Table 5-5 on page 5-47.

Normal Area: Exists within an OSPF domain and connects to the backbone area through one or more ABRs (either physically or through a virtual link). Supports summary link-state advertisements and external link-state advertisements to and from the backbone area, as well as ASBRs.

NSSA (Not-So-Stubby-Area): An OSPF area that limits advertisement of external and summary routes to the backbone area and allows controls on advertisements entering the area from the backbone. (Refer to “Not-So-Stubby-Area (NSSA)” on page 5-53.)

Stub Area: An OSPF area that does not allow an internal ASBR or external type-5 LSAs. (Refer to “Stub Area” on page 5-54.)

Summary Link-State Advertisement: A type-3 LSA summarizing the available links within an OSPF area. This advertisement is sent by the ABR for an area to the backbone area for distribution to the other areas in the OSPF domain. Refer to Table 5-5 on page 5-47. (See also “Link State Advertisement”.)

Type-3 LSA: See “Summary Link-State Advertisement”.

Type-5 LSA: See “External Type-5 Link State Advertisement”.

Type-7 LSA: See “External Type-7 Link State Advertisement”.

Topological Database: See “Link State Database”.

Virtual Link: Used to provide connectivity from a normal area to the backbone when the subject area does not have an ABR physically linked to the backbone area. Refer to “13. Configuring an ABR To Use a Virtual Link to the Backbone” on page 5-89.

Overview of OSPF

OSPF is a link-state routing protocol applied to routers grouped into OSPF areas identified by the routing configuration on each routing switch. The protocol uses link-state advertisements (LSAs) transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces. Each routing switch in an area also maintains a Link State Database (LSDB) that describes the area topology. (All routers in a given OSPF area have identical LSDBs.) The routing switches used to connect areas to each other flood summary link LSAs and external link LSAs to neighboring OSPF areas to update them regarding available routes. Through this means, each OSPF router determines the shortest path between itself and a desired destination router in the same OSPF domain (Autonomous System). Routed traffic in an OSPF AS is classified as one of the following:

- intra-area traffic
- inter-area traffic
- external traffic

The switches covered in this guide support the following types of LSAs, which are described in RFCs 2328 and 3101:

Table 5-5. OSPF LSA Types

LSA Type	LSA Name	Use
1	Router Link	Describes the state of each interface on a router for a given area. Not propagated to backbone area.
2	Network Link	Describes the OSPF routers in a given network. Not propagated to backbone area.
3	Summary Link	Describes the route to networks in another OSPF area of the same Autonomous System (AS). Propagated through backbone area to other areas.
4	Autonomous System (AS) Summary Link	Describes the route to an ASBR in an OSPF Normal or Backbone area of the same AS. Propagated through backbone area to other areas.
5	AS External Link	Describes the route to a destination in another AS (external route). Originated by ASBR in normal or backbone areas of an AS and propagates through backbone area to other normal areas. For injection into an NSSA, ABR converts type-5 LSAs to a type-7 LSA advertising the default route (0.0.0.0/0).
7	AS External Link in an NSSA Area	Describes the route to a destination in another AS (external route). Originated by ASBR in NSSA. ABR converts type-7 LSAs to type-5 LSAs for injection into the backbone area.

OSPF Router Types

Interior Routers

This type of OSPF router belongs to only one area. Interior routers flood type-1 LSAs to all routers in the same area, and maintain identical link state databases (LSDBs). In figure 5-17, below, routers R1, R3, R4, and R6 are all interior routers because all of their links are to other routers in the same area.

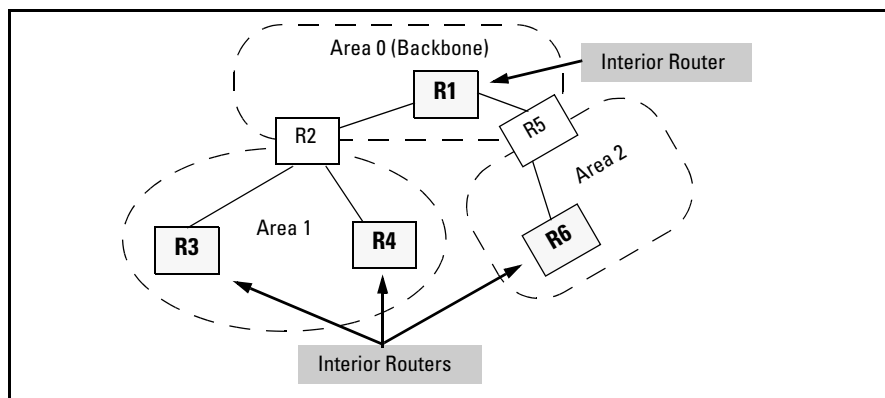


Figure 5-17. Example of Interior Routers

Area Border Routers (ABRs)

This type of OSPF router has membership in multiple areas. ABRs are used to connect the various areas in an AS to the backbone area for that AS. Multiple ABRs can be used to connect a given area to the backbone, and a given ABR can belong to multiple areas other than the backbone. An ABR maintains a separate LSDB for each area to which it belongs. (All routers within the same area have identical LSDBs.) The ABR is responsible for flooding summary LSAs between its border areas. You can reduce summary LSA flooding by configuring area ranges. An area range enables you to assign an aggregate address to a range of IP addresses. This aggregate address is advertised instead of all the individual addresses it represents. You can assign up to eight ranges in an OSPF area. In figure 5-18, below, routers R2 and R5 are Area Border Routers (ABRs) because they both have membership in more than one area.

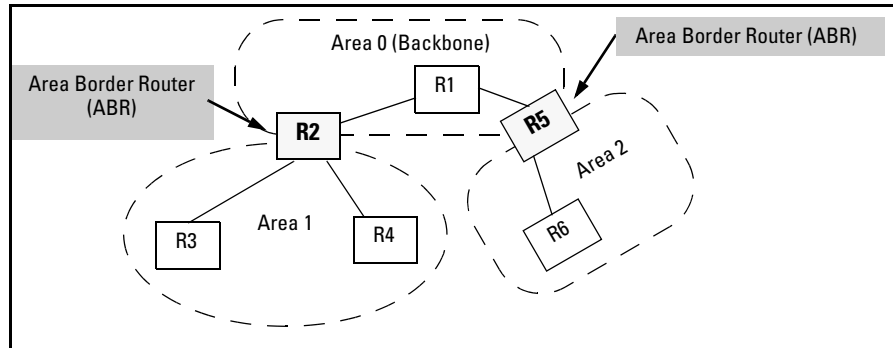


Figure 5-18. Example of Deploying ABRs To Connect Areas to the Backbone

Autonomous System Boundary Router (ASBR)

This type of OSPF router runs multiple Interior Gateway protocols and serves as a gateway to other autonomous systems operating with interior gateway protocols. The ASBR imports and translates different protocol routes into OSPF through *redistribution*. ASBRs can be used in backbone areas, normal areas, and NSSAs, but not in stub areas. For more details on redistribution and configuration examples, see “2. Enable Route Redistribution” on page 5-76.

Designated Routers

In an OSPF network having two or more routers, one router is elected to serve as the designated router (DR) and another router to act as the backup designated router (BDR). All other routers in the area forward their routing information to the DR and BDR, and the DR forwards this information to all of the routers in the network. This minimizes the amount of repetitive information that is forwarded on the network by eliminating the need for each individual router in the area to forward its routing information to all other routers in the network. If the area includes multiple networks, then each network elects its own DR and BDR.

In an OSPF network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next highest priority is elected as the BDR. If the DR goes off-line, the BDR automatically becomes the DR, and the router with the next highest priority then becomes the new BDR. If multiple ProCurve routing switches on the same OSPF network are declaring themselves as DRs, then both priority and router ID are used to select the designated router and backup designated routers.

Priority is configurable by using the **vlan < vid > ip ospf priority < 0-255 >** command at the interface level. You can use this parameter to help bias one router as the DR. (For more on this command, refer to “Priority Per-Interface” on page 5-84.) If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR.

For example, in figure 5-19, the DR and BDR for 10.10.10.0 network in area 5 are determined as follows:

- Router A Priority: 0 Cannot become a DR or BDR.
- Router B Priority: 1 DR for the 10.10.10.0 network.
- Router C Priority: 2 BDR for the 10.10.10.0 network.
- Router D Priority: 0 Cannot become a DR or BDR.
- Router E Priority: 3 Becomes the new BDR if router B becomes unavailable and router C becomes the new DR.




Figure 5-19. Example of Designated Routers in an OSPF Area

To learn the router priority on an interface, use the **show ip ospf interface** command and check the **Pri** setting under **OSPF interface configuration**.

Notes

By default, the router ID is typically the lowest-numbered IP address or the lowest-numbered (user-configured) loopback interface configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 5-15.

If multiple networks exist in the same OSPF area, the recommended approach is to ensure that each network uses a different router as its DR. Otherwise, if a router is a DR for more than one network, latency in the router could increase due to the increased traffic load resulting from multiple DR assignments.

When only one router on an OSPF network claims the DR role despite neighboring routers with higher priorities or router IDs, this router remains the DR. This is also true for BDRs.

The DR and BDR election process is performed when one of the following events occurs:

- an interface is in a waiting state and the wait time expires
- an interface is in a waiting state and a hello packet is received that addresses the BDR
- a change in the neighbor state occurs, such as:
 - a neighbor state transitions from 2 or higher
 - communication to a neighbor is lost
 - a neighbor declares itself to be the DR or BDR for the first time

OSPF Area Types

OSPF is built upon a hierarchy of network areas. All areas for a given OSPF domain reside in the same *Autonomous System (AS)*. An AS is defined as a number of contiguous networks, all of which share the same interior gateway routing protocol.

An AS can be divided into multiple areas. Each area represents a collection of contiguous networks and hosts, and the topology of a given area is not known by the internal routers in any other area. Areas define the boundaries to which types 1 and 2 LSAs are broadcast, which limits the amount of LSA flooding that occurs within the AS and also helps to control the size of the link-state databases (LSDBs) maintained in OSPF routers. An area is represented in OSPF by either an IP address or a number. Area types include:

- backbone
- not-so-stubby (NSSA)
- normal
- stub

All areas in an AS must connect with the backbone through one or more area border routers (ABRs). If a *normal* area is not directly connected to the backbone area, it must be configured with a *virtual link* to an ABR that is directly connected to the backbone. The remaining area types do not allow virtual link connections to the backbone area.

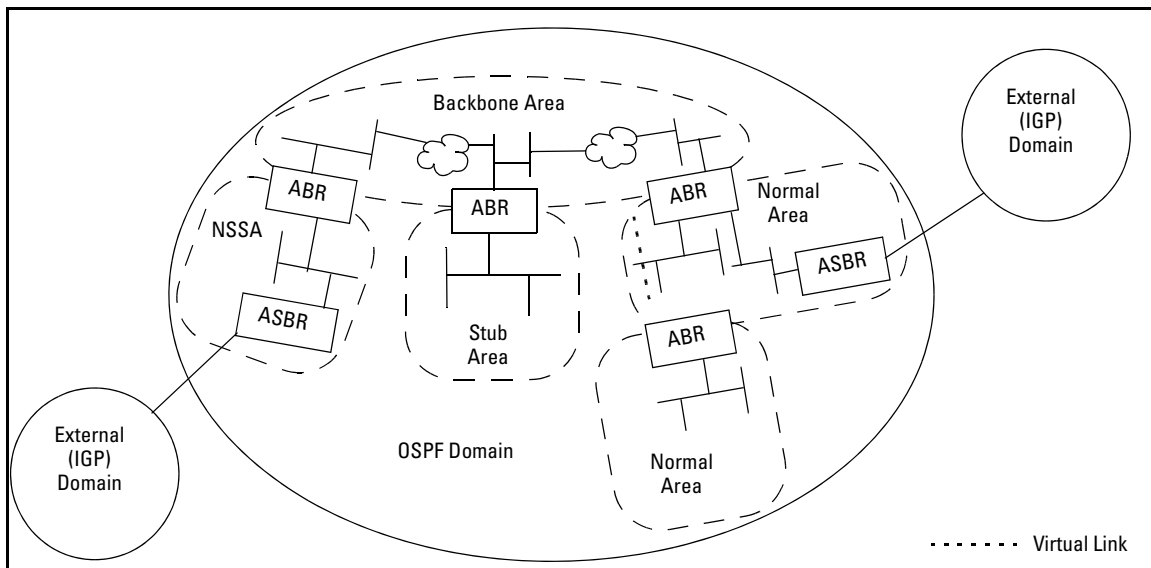


Figure 5-20. Example of an Autonomous System (AS) with Multiple Areas and External Routes

Backbone Area

Every AS must have one (and only one) backbone area (identified as area 0 or 0.0.0.0). The ABRs of all other areas in the same AS connect to the backbone area, either physically through an ABR or through a configured, virtual link. The backbone is a transit area that carries the type-3 summary LSAs, type-5 AS external link LSAs and routed traffic between non-backbone areas, as well as the type-1 and type-2 LSAs and routed traffic internal to the area. ASBRs are allowed in backbone areas.

Normal Area

This area connects to the AS backbone area through one or more ABRs (physically or through a virtual link) and supports type-3 summary LSAs and type-5 external link LSAs to and from the backbone area. ASBRs are allowed in normal areas.

Not-So-Stubby-Area (NSSA)

Beginning with software release K.12.33, this area is available and connects to the backbone area through one or more ABRs. NSSAs are intended for use where an ASBR exists in an area where you want to control the following:

- advertising the ASBR's external route paths to the backbone area
- advertising the NSSA's summary routes to the backbone area
- allowing LSAs from the backbone area to advertise in the NSSA:
 - summary routes (type-3 LSAs) from other areas
 - external routes (type-5 LSAs) from other areas as a default external route (type-7 LSAs)

In the above operation, the ASBR in the NSSA injects external routes as type 7 LSAs. (Type 5 LSAs are not allowed in an NSSA.) The ABR connecting the NSSA to the backbone converts the type 7 LSAs to type 5 LSAs and injects them into the backbone area for propagation to networks in the backbone and to any normal areas configured in the AS. The ABR also injects type-3 summary LSAs:

- from the NSSA into the backbone area
- from the backbone into the NSSA

As mentioned above, if the ABR detects type-5 external LSAs on the backbone, it injects a corresponding type-7 LSA default route (0.0.0.0/0) into the NSSA

You can also configure the NSSA ABR to do the following:

- Suppress advertising some or all of the area's summarized internal or external routes into the backbone area. (Refer to "8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone" on page 5-77.)
- Replace all type-3 summary routes and the type-7 default route with the type-3 default summary route (0.0.0.0/0).

Virtual links are not allowed for NSSAs.

Stub Area

This area connects to the AS backbone through one or more ABRs. It does not allow an internal ASBR, and does not allow external (type 5) LSAs. A stub area supports these actions:

- Advertise the area's summary routes to the backbone area.
- Advertise summary routes from other areas.
- Use the default summary (type-3) route to advertise both of the following:
 - summary routes to other areas in the AS
 - external routes to other autonomous systems

You can configure the stub area ABR to do the following:

- Suppress advertising some or all of the area's summarized internal routes into the backbone area.
- Suppress LSA traffic from other areas in the AS by replacing type-3 summary LSAs and the default external route from the backbone area with the default summary route (0.0.0.0/0).

Virtual links are not allowed for stub areas.

OSPF RFC Compliance

The OSPF features covered in this guide comply with the following:

- RFC 2328 OSPF version 2
- RFC 3101 OSPF NSSA option (s/w release K.12.*xx* and greater)
- RFC 1583 (Enabled in the default OSPF configuration. Refer to the following Note.)

Note

If all of the routers in your OSPF domain support RFC 2178, RFC 2328, or later, you should disable RFC 1583 compatibility on all routers in the domain. Refer to “3. Changing the RFC 1583 OSPF Compliance Setting” on page 5-64.

Reducing AS External LSAs and Type-3 Summary LSAs

An OSPF ASBR uses AS External LSAs to originate advertisements of a route to another routing domain, such as a RIP domain. These advertisements are

- flooded in the area in which the ASBR operates

- injected into the backbone area and then propagated to any other OSPF areas (except stub areas) within the local OSPF Autonomous System (AS). If the AS includes an NSSA, there are two additional options:
 - If the NSSA includes an ASBR, you can suppress advertising some or all of its summarized external routes into the backbone area.
 - Replace all type-3 summary LSAs and the default external route from the backbone area with the default summary route (0.0.0.0/0).

In some cases, multiple ASBRs in an AS can originate equivalent external LSAs. The LSAs are equivalent when they have the same cost, the same next hop, and the same destination. In such cases, the ProCurve switch optimizes OSPF by eliminating duplicate AS External LSAs. That is, the ASBR with the highest router ID floods the AS External LSAs for the external domain into the OSPF AS, while the other ASBRs flush the equivalent AS External LSAs from their databases. As a result, the overall volume of route advertisement traffic within the AS is reduced and the switches that flush the duplicate AS External LSAs have more memory for other OSPF data.

This enhancement implements the portion of RFC 2328 that describes AS External LSA reduction. This enhancement is enabled by default, requires no configuration, and cannot be disabled.

Algorithm for AS External LSA Reduction

The AS External LSA reduction feature behavior changes under the following conditions:

- There is one ASBR advertising (originating) a route to the external destination, but one of the following happens:
 - A second ASBR comes on-line
 - A second ASBR that is already on-line begins advertising an equivalent route to the same destination.

In either case above, the ProCurve switch with the higher router ID floods the AS External LSAs and the other ProCurve switch flushes its equivalent AS External LSAs.

- One of the ASBRs starts advertising a route that is no longer equivalent to the route the other ASBR is advertising. In this case, the ASBRs each flood AS External LSAs. Since the LSAs either no longer have the same cost or no longer have the same next-hop router, the LSAs are no longer equivalent, and the LSA reduction feature no longer applies.
- The ASBR with the higher router ID becomes unavailable or is reconfigured so that it is no longer an ASBR. In this case, the other ASBR floods the AS External LSAs.

Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA

By default, a routing switch operating as an ABR for a stub area or NSSA injects non-default, summary routes (LSA type 3) into the stub areas and NSSAs. For NSSAs, the routing switch also injects a type-7 default external route. You can further reduce LSA traffic into these areas by using **no-summary**. This command option configures the routing switch to:

- Replace type-3 summary LSA injection into a stub area or NSSA with a type-3 default summary route (0.0.0.0/0).
- Disable injection of the type-7 default external route into an NSSA.

You can enable this behavior when you first configure the stub area or NSSA, or at a later time. (For the full command to use, refer to “Configuring a Stub or NSSA Area” on page 5-68.)

The **no-summary** command does not affect intra-area advertisements, meaning the switch still accepts summary LSAs from OSPF neighbors within its area and floods them to other neighbors. The switch can form adjacencies with other routers regardless of whether summarization is enabled or disabled for areas on each switch.

When you use **no-summary**, the change takes effect immediately. If you apply the option to a previously configured area, the switch flushes all of the summary LSAs it has generated (as an ABR) from the area.

Note

This feature applies only when the switch is configured as an Area Border Router (ABR) for a stub area or NSSA. To completely prevent summary LSAs from injection into the area, use **no-summary** to disable the summary LSAs on each OSPF router that is an ABR for the area.

To implement the above operation for a stub area or NSSA, enter a command such as the following:

```
ProCurve (ospf) # area 40 stub 3 no-summary
```


Equal Cost Multi-Path Routing

The Equal Cost Multi-Path (ECMP) feature allows OSPF to add routes with multiple next-hop addresses and with equal costs to a given destination in the Forwarding Information Base (FIB) on the routing switch. For example, if you display the IP Route table by entering the **show ip route** command, multiple next-hop routers are listed for the same destination network (21.0.9.0/24) as shown in Figure 5-21.

```
ProCurve> show ip route
```

IP Route Entries						
Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist.
1.0.0.0/8	10.0.8.1	1	static		1	1
10.0.8.0/21	DEFAULT_VLAN	1	connected		1	0
12.0.9.0/24	VLAN3	3	connected		1	0
15.0.0.0/8	10.0.8.1	1	static		1	1
21.0.9.0/24	162.130.101.2	2	ospf	IntraArea	2	110
21.0.9.0/24	162.130.101.3	2	ospf	IntraArea	2	110
21.0.9.0/24	162.130.101.4	2	ospf	IntraArea	2	110
127.0.0.0/8	reject		static		0	0
127.0.0.1/32	lo0		connected		1	0
162.130.101.0/24	VLAN2	2	connected		1	0

Multiple next-hop gateway addresses are displayed for the destination network 21.0.9.0/24.

Figure 5-21. "Example of "show ip route" Command Output with Multiple Next-Hop Routes

For a given destination network in an OSPF domain, multiple ECMP next-hop routes can be *one* of the following types.

- Intra-area (routes to the destination in the same OSPF area)
- Inter-area (routes to the destination through another OSPF area)
- External (routes to the destination through another autonomous system)

Multiple ECMP next-hop routes cannot be a mixture of intra-area, inter-area, and external routes. For example, in Figure 5-21, the multiple next-hop routes to network 21.0.9.0/24 are all intra-area.

Also, according to the distributed algorithm used in the selection of ECMP next-hop routes:

- Intra-area routes are preferred to inter-area routes.
- Inter-area routes are preferred to external routes through a neighboring autonomous system.

In addition, ECMP ensures that all traffic forwarded to a given host address follows the same path, which is selected from the possible next-hop routes.

For example, in Figure 5-22, the ECMP inter-area routes to destination network 10.10.10.0/24 consist of the following next-hop gateway addresses: 12.0.9.2, 13.0.9.3, and 14.0.9.4.

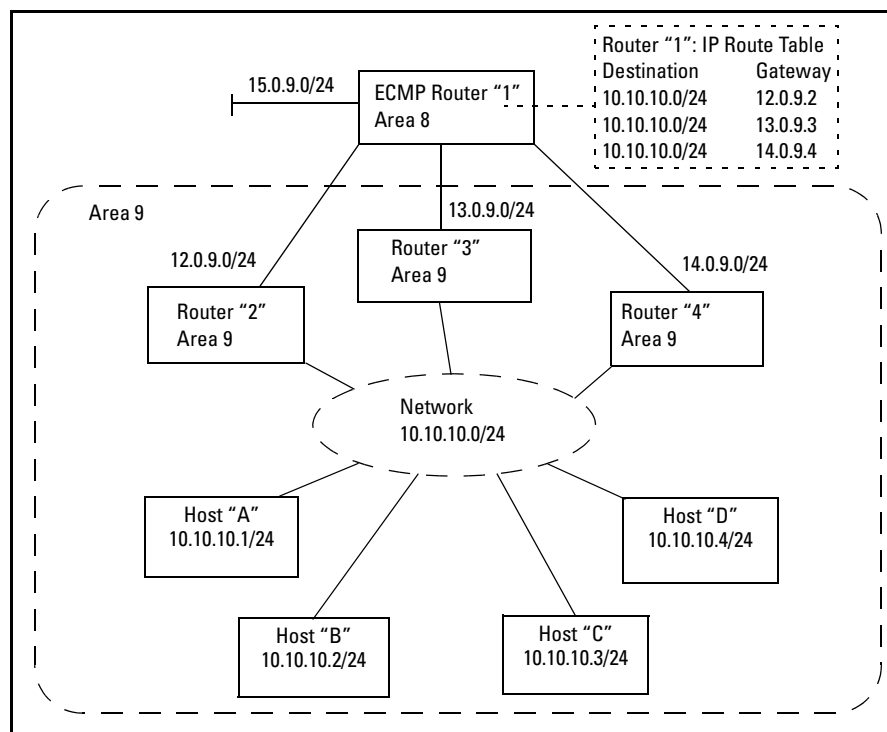


Figure 5-22. Example of OSPF ECMP Multiple Next-Hop Routing (Inter-Area)

However, the forwarding software distributes traffic across the three possible next-hop routes in such a way that all traffic for a specific host is sent to the same next-hop router.

As shown in Figure 5-23, one possible distribution of traffic to host devices is as follows:

- Traffic to host 10.10.0.1 passes through next-hop router 12.0.9.2.
- Traffic to host 10.10.0.2 passes through next-hop router 13.0.9.3.
- Traffic to host 10.10.0.3 passes through next-hop router 12.0.9.2.
- Traffic to host 10.10.0.4 passes through next-hop router 14.0.9.4.

IP Packet Destination	Next Hop Used
10.10.0.1	12.0.9.2
10.10.0.2	13.0.9.3
10.10.0.3	12.0.9.2
10.10.0.4	14.0.9.4

Figure 5-23. Example of Traffic Distribution on ECMP Next-Hop Routers

Dynamic OSPF Activation and Configuration

OSPF automatically activates when enabled with **router ospf**. All configuration commands affecting OSPF (except reconfiguring the router ID) are dynamically implemented, and can be used without restarting OSPF routing. (To reconfigure the router ID, refer to “Changing the Router ID” on page 5-15.)

Note

OSPF is automatically enabled without a system reset.

General Configuration Steps for OSPF

To begin using OSPF on the routing switch, perform the steps outlined below:

1. In the global config context, use **ip routing** to enable routing (page 5-63).
2. Execute **router ospf** to place the routing switch in the **ospf** context and to enable OSPF routing (page 5-63).
3. Change the OSPF RFC 1583 compliance, if needed. (Refer to “3. Changing the RFC 1583 OSPF Compliance Setting” on page 5-64.)
4. Use **area** to assign the areas to which the routing switch will be attached (page 5-66).
5. Assign interfaces to the configured areas per-VLAN or per-subnet by moving to each VLAN context and using one of the following commands:
 - **ip ospf area < ospf-area-id >** assigns all interfaces in the VLAN to the same area. Use this option when there is only one IP address configured on the VLAN or you want all subnets in the VLAN to belong to the same OSPF area.
 - **ip ospf < ip-address > area < ospf-area-id >** assigns an individual subnet to the specified area.

(Refer to page 5-71.)

6. Optional: Assign loopback interfaces to OSPF areas by using the **ip ospf area** command at the loopback interface configuration level.
(Refer to page 5-72.)
7. Optional: On each routing switch used as an ASBR in your OSPF domain, configure redistribution to enable importing the routes you want to make available in the domain.
 - i. On an ASBR in a backbone, normal, or NSSA area where you want to import external routes, configure redistribution filters to define the external routes you *do not* want imported.
 - ii. Enable redistribution.

Refer to “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-74.

8. Optional: Configure ranges on ABRs to reduce inter-area route advertising.
9. Optional: Use Administrative Distance to influence route choices.
10. Optional: Change OSPF trap generation.
11. Optional: Reconfigure default parameters in the interface context, if needed. Includes **cost**, **dead-interval**, **hello-interval**, **priority**, and others.

12. Optional: Configure OSPF interface authentication.
13. Configure virtual links for any areas not directly connected to the backbone.

Configuration Rules

- If the switch is to operate as an ASBR, you must enable redistribution (step 7 on page 5-60). When you do that, ASBR capability is automatically enabled. For this reason, you should first configure redistribution filters on the ASBR. Otherwise, all possible external routes will be allowed to flood the domain. (Refer to “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-74.)
- Each VLAN interface on which you want OSPF to run must be assigned to one of the defined areas. When a VLAN interface is assigned to an area, the IP address is automatically included in the assignment. To include additional addresses, you must enable OSPF on them separately, or use the “all” option in the assignment.

OSPF Global and Interface Settings

When first enabling OSPF, you may want to consider configuring ranges and restricting redistribution (if an ASBR is used) to avoid unwanted advertisements of external routes. You may also want to enable the OSPF trap and authentication features to enhance troubleshooting and security. However, it is generally recommended that the remaining parameters with non-null default settings be left as-is until you have the opportunity to assess OSPF operation and determine whether any adjustments to non-default settings is warranted.

The following tables list the global and per-interface commands used with OSPF. For information on when to use these commands, refer to “General Configuration Steps for OSPF” on page 5-60. For detailed information on each command, refer to the page listed for each command.

Table 5-6. OSPF Default Global Settings

Parameter	Default	Page
area < area-#> virtual-link < ip-addr >	None	
default-metric	10	
distance < external inter-area intra-area >	110	
range	All IP Addresses	
redistribute	Disabled	

Parameter	Default	Page
restrict	Disabled	
rfc-1583-compatibility	Enabled	
metric-type	type2	
trap < ospf-trap >	Enabled	

Note

Set global level parameters in the **ospf** context of the CLI. To access this context level, ensure that routing is enabled, then execute **router ospf** at the global CONFIG level. For example:

```
ProCurve (config)# router ospf
ProCurve (ospf)#
```

Table 5-7. OSPF Default Interface Settings

Parameter	Default	Page
all	n/a	
area	None	
authentication-key	None	
cost	1	
dead-interval	40 seconds	
hello-interval	10 seconds	
IP-ADDR	None	
md5-auth-key-chain	None	
priority	1	
retransmit-interval	5 seconds	
transit-delay	1 second	

Note

Use the VLAN interface context to set interface level OSPF parameters for the desired VLAN. To access this context level, use **vlan < vid >** either to move to the VLAN context level or to specify that context from the global config level. For example, both of the following two command sets achieve the same result:

```
ProCurve (config)# vlan 20
ProCurve (vlan-20)# cost 15

ProCurve (config)# vlan 20 cost 15
```

Configuring OSPF on the Routing Switch

1. Enable IP Routing

Syntax: [no] ip routing

Executed at the global configuration level to enable IP routing on the routing switch.

Default: Disabled

*The **no** form of the command disables IP routing. (Global OSPF and RIP routing must be disabled before you disable IP routing.)*

```
ProCurve(config)# ip routing
```

2. Enable Global OSPF Routing

Syntax: [no] router ospf

Executed at the global configuration level to enable OSPF on the routing switch and to enter the OSPF router context. This enables you to proceed with assigning OSPF areas, including ABR and ASBR configuration, and to modify OSPF global parameter settings as needed. Global IP routing must be enabled before executing this command.

Default: Disabled

*The **no** form of the command disables OSPF routing.*

Note: *If you disable OSPF, the switch retains all the configuration information for the disabled protocol in flash memory. If you subsequently restart OSPF, the existing configuration will be applied.*

For example:

```
ProCurve(config)#router ospf
ProCurve(ospf)#
```

3. Changing the RFC 1583 OSPF Compliance Setting

In OSPF domains supporting multiple external routes from different areas to the same external destination, multiple AS-external-LSAs advertising the same destination are likely to occur. This can cause routing loops and the network problems that loops typically generate. On the routing switches covered by this guide, if RFC 1583 compatibility is disabled, the preference rules affecting external routes are those stated in RFC-2328, which minimize the possibility of routing loops when AS-external-LSAs for the same destination originate from ASBRs in different areas. However, because all routers in an OSPF domain must support the same routing-loop prevention measures, if the domain includes any routers that support only RFC 1583 preference rules, then all routers in the domain must be configured to support RFC 1583.

Note

The routing switch is configured, by default, to be compliant with the RFC 1583 OSPF V2 specification. (Use **show ip ospf general** to view the current RFC 1583 configuration setting.)

All routes in an AS should be configured with the same compliance setting for preference rules affecting external routes. Thus, if any routers in an OSPF domain support only RFC 1583, then all routers must be configured with 1583 compatibility. In the default OSPF configuration, RFC 1583 support is enabled for the routing switches covered by this guide.

If all routers in the domain support RFC 2178 or RFC 2328, then you should disable RFC 1583 compatibility on all of the routers, since conformance to these later RFCs provides more robust protection against routing loops on external routes.

Syntax: [no] rfc1583-compatibility

Executed at the global configuration level to toggle routing switch operation compliance between RFC 1583 and RFC 2328.

rfc1583-compatibility: *Configures the routing switch for external route preference rules compliant with RFC 1583.*

no rfc1583-compatibility: *Configures the routing switch for external route preference rules compliant with RFC 2328.*

Default: Compliance enabled

For example, to disable RFC 1583 compatibility on a routing switch in an OSPF domain where RFC 2178 and RFC 2328 are universally supported:

```
ProCurve(config)# router ospf
ProCurve(ospf)# no rfc1583-compatibility
```



```
ProCurve(config)# router ospf
ProCurve(ospf)# no rfc1583-compatibility
ProCurve_8212(ospf)# show ip ospf general
```

OSPF General Status

OSPF protocol	: enabled
Router ID	: 10.10.51.1
RFC 1583 compatibility	: non-compatible
Intra-area distance	: 110
Inter-area distance	: 110
AS-external distance	: 110
Default import metric	: 10
Default import metric type	: external type 2
Area Border	: no
AS Border	: yes
External LSA Count	: 9
External LSA Checksum Sum	: 408218
Originate New LSA Count	: 24814
Receive New LSA Count	: 14889

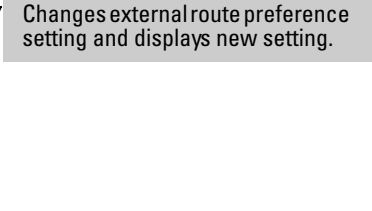


Figure 5-24. Example of Changing External Route Preference Compatibility from RFC 1583 to RFC 2328

4. Assign the Routing Switch to OSPF Areas

After you globally enable OSPF on the routing switch (in the previous step), use this command to assign one or more OSPF areas within your autonomous system (AS). A routing switch can belong to one area or to multiple areas. (Participation in a given, assigned area requires configuring one or more VLANs or subnets and assigning each to the desired area. Refer to page 5-71.)

- If you want the VLANs and any subnets configured on the routing switch to all reside in the same area, then you need to configure only that one area. (In this case, the routing switch would operate as an internal router for the area.)
- If you want to put different VLANs or subnets on the routing switch into different areas, then you need to re-execute this command for each area. (In this case, the routing switch will operate as an ABR for each of the configured areas.)

Note

Each ABR must either be directly connected to the backbone area (0) or be configured with a virtual link to the backbone area through another ABR that is directly connected to the backbone area. For information on this step, refer to “13. Configuring an ABR To Use a Virtual Link to the Backbone” on page 5-89.

Configuring an OSPF Backbone or Normal Area.

Syntax: area < ospf-area-id | backbone > [normal]
no area < ospf-area-id | backbone >

*After using **router ospf** to globally enable OSPF and enter the global OSPF context, execute this command to assign the routing switch to a backbone or other normal area.*

*The **no** form of the command removes the routing switch from the specified area.*

Default: No areas. Range: 1-16 areas (of all types)

*< ospf-area-id >: Specifies a normal area to which you are assigning the routing switch. You can assign the routing switch to one or more areas, depending on the area in which you want each configured VLAN or subnet to reside. You can enter area IDs in either whole number or dotted decimal format. (The routing switch automatically converts whole numbers to the dotted decimal format.) For example, if you enter an area-ID of **1**, it appears in the switch's configuration as **0.0.0.1** and an area-ID of 256 appears in the switch configuration as **0.0.1.0**. An area ID can be a value selected to match the IP address of a VLAN belonging to the area, or a value corresponding to a numbering system you devise for the areas in a given AS. Entering an area ID of **0** or **0.0.0.0** automatically joins the routing switch to the Backbone area. The maximum area ID value is 255.255.255.254 (4,294,967,294).*

backbone: *Assigns the routing switch to the backbone area and automatically assigns an area ID of **0.0.0.0** and an area type of **normal**. Using **0** or **0.0.0.0** with the above **ospf-area-id** option achieves the same result. The backbone area is automatically configured as a "normal" area type.*

For example, to configure a backbone and a normal area with an ID of "1" (0.0.0.1) on a routing switch:

```
ProCurve (ospf) # area backbone
ProCurve (ospf) # area 1
```

Configuring a Stub or NSSA Area.

Syntax: area < ospf-area-id > stub < 0-16777215 > [no-summary]
area < ospf-area-id > nssa < 0-16777215 > [no-summary]
[metric-type < type1 | type2 >]
no area < ospf-area-id >

*After using **router ospf** to globally enable OSPF and enter the global OSPF context, execute this command to assign the routing switch to a stub area or NSSA. (Does not apply to backbone and normal OSPF area ABRs.)*

*The **no** form of the command removes the routing switch from the specified area.*

Default: No areas. Range: 1-16 areas (of all types)

< ospf-area-id >: Same area ID as on page 5-67 except you cannot assign a backbone area number (**0** or **0.0.0.0**) to a stub or NSSA area.

< stub | nssa > Designates the area identified by **< ospf-area-id >** as a stub area or NSSA.

< 0-16777215 >: If the routing switch is used as an ABR for the designated area, assigns the cost of the default route (to the backbone) that is injected into the area.

Notes: If the routing switch is not an ABR for the stub area or NSSA, the above cost setting is still required by the CLI, but is not used.

In the default configuration, a routing switch acting as an ABR for a stub area or NSSA injects type-3 summary routes into the area. For an NSSA, the routing switch also injects a type-7 default route into the area.

[no-summary]: Where the routing switch is an ABR for a stub area or an NSSA, this option reduces the amount of LSA traffic entering the area from the backbone by replacing the injection of type-3 summary routes with injection of a type-3 default summary route. For NSSAs, this command also disables injection of the type-7 default external route from the backbone into the area (included in the **metric-type** operation described below).

(Default: Disabled)

For more on this topic, refer to “Not-So-Stubby-Area (NSSA)” on page 5-53, “Stub Area” on page 5-54, and “Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA” on page 5-56.

[metric-type < type1 | type2 >]: Used in NSSA ABRs only. Enables injection of the type-7 default external route and type-3 summary routes into the area instead of a type 3 default route. Also specifies the type of internal cost metric to include in type-7 LSAs advertised for redistribution of external routes in the NSSA. (The redistribution—or external—cost metric is a global setting on the routing switch set by the **default-metric** command.) The **metric-type** command specifies whether to include the redistribution cost in the cost metric calculation for a type-7 default LSA injected into the area.

type1: Calculate external route cost for a type-7 default LSA as the sum of (1) the external route cost assigned by the ASBR plus (2) the internal cost from the router with traffic for the external route to the ASBR advertising the route.

type2: Calculate external route cost for a type-7 default LSA as being only the cost from the router with traffic for the external route to the ASBR advertising the route.

Using the **area < ospf-area-id > nssa < 0-16777215 >** without entering either **no-summary** or **metric-type** resets the routing switch to the state where injection of type-3 summary routes and the type-7 default external routes is enabled with **metric-type** set to **type2**.

(Default: Enabled with **metric-type type2**.)

Note: Different routers in the NSSA can be configured with different **metric-type** values.

The following examples of configuring a stub area and an NSSA on a routing switch use an (arbitrary) cost of “10”.

```
ProCurve(ospf)# area 2 stub 10
ProCurve(ospf)# area 3 nssa 10
ProCurve(ospf)# area 4 nssa 10 no-summary
ProCurve(ospf)# area 5 nssa 10 metric-type type1
```

Assigns a stub area with a cost of 10.

Assigns an NSSA with a cost of 10 and, by default, uses a Type2 default cost metric for Type-7 (external) route LSAs received from the backbone.

Assigns an NSSA with a cost of 10, blocks injection of type-3 summary routes, and starts injection of type-3 default routes from the backbone.

Sets the cost metric type for type-7 (default) LSAs injected into the NSSA.

Figure 5-25. Examples of Creating Stub Area and NSSA Assignments

5. Assign VLANs and/or Subnets to Each Area

After you define an OSPF area (page 5-66), you can assign one or more VLANs and/or subnets to it. When a VLAN is assigned to an area, all currently configured IP addresses in that VLAN are automatically included in the assignment unless you enter a specific IP address.

Note

All static VLANs configured on a routing switch configured for OSPF must be assigned to one of the defined areas in the AS.

Syntax: `vlan < vid ># ip ospf [ip-addr | all] area < ospf-area-id >`

Executed in a specific VLAN context to assign the VLAN or individual subnets in the VLAN to the specified area. Requires that the area is already configured on the routing switch (page 5-66). When executed without specifying an IP address or using the **all** keyword, this command assigns all configured networks in the VLAN to the specified OSPF area.

vlan < vid >: Defines the VLAN context for executing the area assignment.

[ip-addr]: Defines a specific subnet on the VLAN to assign to a configured OSPF area.

[all]: Assigns all subnets configured on the VLAN to a configured OSPF area.

area < ospf-area-id >: Identifies the OSPF area to which the VLAN or selected subnet should be assigned.

Notes: *If you add a new subnet IP address to a VLAN after assigning the VLAN to an OSPF area, you must also assign the new subnet to an area. If all subnets in the VLAN should be assigned to the same area, just execute **ip ospf area < ospf-area-id >**. But if different subnets belong in different areas, you must explicitly assign the new subnet to the desired area. Also, to assign a VLAN to an OSPF area, the VLAN must be configured with at least one IP address. Otherwise, executing this command results in the following CLI message:*

OSPF can not be configured on this VLAN.

Example: To assign VLAN 8 on a routing switch to area 3 and include *all* IP addresses configured in the VLAN, enter the following commands:

```
ProCurve(ospf)# vlan 8
ProCurve(vlan-8)# ip ospf area 3
```

Example. Suppose that a system operator wants to assign the three subnets configured in VLAN 10 as shown below:

- 10.10.10.1 to OSPF area 5
- 10.10.11.1 to OSPF area 5
- 10.10.12.1 to OSPF area 6

The operator could use the following commands to configure the above assignments:

```
ProCurve (ospf) # vlan 10
ProCurve (vlan-10) # ip ospf 10.10.10.1 area 5
ProCurve (vlan-10) # ip ospf 10.10.11.1 area 5
ProCurve (vlan-10) # ip ospf 10.10.12.1 area 6
```

6. Optional: Assigning Loopback Addresses to an Area

After you define the OSPF areas to which the switch belongs, you can assign a user-defined loopback address to an OSPF area. A loopback interface is a virtual interface configured with an IP address and is always reachable as long as at least one of the IP interfaces on the switch is operational. Because the loopback interface is always up, you ensure that the switch's router ID remains constant and that an OSPF network is protected from changes caused by downed interfaces.

For more information about how to configure a loopback interface, refer to “Configuring a Loopback Interface” in the chapter titled, “Configuring IP Addressing”, in the *Management and Configuration Guide* for your routing switch.

Syntax: interface loopback <0-7> ip ospf <lo-ip-address> area <ospf-area-id>

Executed in a specific loopback context to assign a loopback interface to the specified OSPF area. Requires that the specified loopback interface is already configured with an IP address on the switch.

loopback interface <0-7>: Defines the loopback context for executing the area assignment.

ip ospf <lo-ip-address>: Specifies the loopback interface by its IP address to assign to a configured OSPF area.

area <ospf-area-id>: Identifies the OSPF area to which the loopback interface is assigned. You can enter a value for the OSPF area in the format of an IP address or a number in the range 0 to 4,294,967,295.

Example: To assign user-defined loopback interface 3 on the switch to area 192.5.0.0 and include the loopback IP address 172.16.112.2 in the OSPF broadcast area, enter the following commands:

```
ProCurve(config)# interface loopback 3
ProCurve(lo-3)# ip ospf 172.16.112.2 area 192.5.0.0
```

Syntax: interface loopback <0-7># ip ospf < lo-ip-address > cost < number >

Executed in a specific loopback context to modify the cost used to advertise the loopback address (and subnet) to the area border router (ABR). Requires that the specified loopback interface is already configured with an IP address on the switch.

loopback interface <0-7>: Defines the loopback context for executing the cost assignment.

ip ospf < lo-ip-address >: Specifies the loopback interface by its IP address.

cost < number >: Specifies a number that represents the administrative metric associated with the loopback interface. Valid values are from 1 to 65535. Default: 1.

Example: To configure a cost of 10 for advertising the IP address 172.16.112.2 configured for loopback interface 3 in an OSPF area 192.5.0.0, enter the following commands:

```
ProCurve(config)# interface loopback 3
ProCurve(lo-3)# ip ospf 172.16.112.2 area 192.5.0.0
ProCurve(lo-3)# ip ospf 172.16.112.2 cost 10
```

OSPF Redistribution of Loopback Addresses: When you assign a loopback address to an OSPF area, the route redistribution of the loopback address is limited to the specified area.

When route redistribution is enabled:

- The switch advertises a loopback IP address that is not assigned to an OSPF area as an OSPF *external* route to its OSPF neighbors, and handles it as a connected route.
- The switch advertises a loopback address that is assigned to an OSPF area as an OSPF *internal* route.

To enable redistribution of loopback IP addresses in OSPF, enter the **redistribution connected** command as described in “2. Enable Route Redistribution” on page 5-76.

Example: In the following configuration, the loopback IP address 13.3.4.5 of loopback 2 is advertised only in OSPF area 0.0.0.111. The IP addresses 14.2.3.4 and 15.2.3.4 of loopback 1 are advertised in all OSPF areas.

```
ProCurve(config)# interface loopback 1
ProCurve(lo-1)# ip address 14.2.3.4
ProCurve(lo-1)# ip address 15.2.3.4
ProCurve(lo-1)# exit
ProCurve(config)# interface loopback 2
ProCurve(lo-2)# ip address 13.3.4.5
ProCurve(lo-2)# ip ospf 15.2.3.4 area 0.0.0.111
ProCurve(lo-2)# exit
```

Assigns the IP address of loopback interface 2 to OSPF area 111.

Figure 5-26. Examples of Assigning Loopback IP Addresses to OSPF Areas

To verify the OSPF redistribution of loopback interfaces, enter the **show ip route** command from any context level to display IP route table entries.

Note that in the following example, a loopback address assigned to an area is displayed as an **ospf intra-area** (internal) route to its neighbor; a loopback address not assigned to a specific area is displayed as an **ospf external** route:

```
ProCurve(config)# show ip route
```

IP Route Entries						
Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist
20.0.15.1/32	25.0.67.131	25	ospf	external2	10	110
20.0.16.2/32	25.0.67.131	25	ospf	intra-area	2	110

Figure 5-27. Example of Verifying OSPF Redistribution of Loopback Interfaces

7. Optional: Configure for External Route Redistribution in an OSPF Domain

Configuring route redistribution for OSPF establishes the routing switch as an ASBR (residing in a backbone, normal, or NSSA) for importing and translating different protocol routes from other IGP domains into an OSPF domain. The switches covered by this guide support redistribution for static routes, RIP routes, and directly connected routes from RIP domains into OSPF

domains. When you configure redistribution for OSPF, you can specify that static, connected, or RIP routes external to the OSPF domain are imported as OSPF routes. (Likewise, RIP redistribution supports the import of static, connected, and OSPF routes into RIP routes.) The steps for configuring external route redistribution to support ASBR operation include the following:

1. Configure redistribution filters to exclude external routes that you do not want redistributed in your OSPF domain.
2. Enable route redistribution.
3. Modify the default metric for redistribution (optional).
4. Modify the redistribution metric type (optional).
5. Change the administrative distance setting (optional).

Note

Do not enable redistribution until you have used **restrict** to configure the redistribution filters. Otherwise, your network might get overloaded with routes that you did not intend to redistribute.

1. Configure Redistribution Filters.

Syntax: `router ospf restrict < ip-addr/mask-length >`

This command prevents distribution of the specified range of external routes through an ASBR from sources external to the OSPF domain.

Default: Allow all supported, external route sources.

Note: *Use this command to block unwanted, external routes before enabling route redistribution on the ASBR.*

Example: To configure a routing switch operating as an ASBR to filter out redistribution of static, connected, or RIP routes on network 10.0.0.0, enter the following commands:

```
ProCurve(config)# router ospf restrict 10.0.0.0/8
```

Note

In the default configuration, redistribution is permitted for all routes from supported sources.

2. Enable Route Redistribution. This step enables ASBR operation on a routing switch, and must be executed on each routing switch connected to external routes you want to redistribute in your OSPF domain.

Note

Do not enable redistribution until you have configured the redistribution “restrict” filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Syntax: [no] router ospf redistribute < connected | static | rip >

Executed on an ASBR to globally enable redistribution of the specified route type to the OSPF domain through the area in which the ASBR resides.

static: *Redistribute from manually configured routes.*

connected: *Redistribute from locally connected network(s).*

rip: *Redistribute from RIP routes.*

*The **no** form of the command disables redistribution for the specified route type.*

For example, to enable redistribution of all supported external route types through a given ASBR, execute the following commands.

```
ProCurve(config)# router ospf redistribution connected
ProCurve(config)# router ospf redistribution static
ProCurve(config)# router ospf redistribution rip
```

3. Modify the Default Metric for Redistribution. The default metric is a global parameter that specifies the cost applied to all OSPF routes by default

Syntax: router ospf default-metric < 0-16777215 >

Globally assigns the cost metric to apply to all external routes redistributed by the ASBR. By using different cost metrics for different ASBRs, you can prioritize the ASBRs in your AS.

Default: 10; Range: 0-16777215

Example: To assign a default metric of 4 to all routes imported into OSPF on an ASBR, enter the following commands:

```
ProCurve()#
ProCurve(config)# router ospf default-metric 4
```

4. Modifying the Redistribution Metric Type. The redistribution metric type is used by default for all routes imported into OSPF. Type 1 metrics are the same “units” as internal OSPF metrics and can be compared directly. Type 2 metrics are not directly comparable, and are treated as larger than the largest internal OSPF metric.

Syntax: `router ospf metric-type < type1 | type2 >`

Globally reconfigures the redistribution metric type on an ASBR.

type1: Specifies the OSPF metric plus the external metric for an external route.

type2: Specifies the external metric for an external route.

Default: type2

For example, to change from the default setting on an ASBR to type 1, enter the following command:

```
ProCurve(config)# router ospf metric-type type1
```

8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone

Configuring ranges does the following to reduce inter-area advertising:

- **Summarizing Routes:** Enable a routing switch operating as an ABR to use a specific IP address and mask to summarize a range of IP addresses into a single route advertisement for injection into the backbone. This results in only one address being advertised to the network instead of all the addresses within that range. This reduces LSA traffic and the resources needed to maintain routing tables.
- **Blocking Routes:** Prevent an ABR from advertising specific networks or subnets to the backbone area.

Each OSPF area supports up to 8 range configurations.

Syntax: `area < ospf-area-id > range < ip-addr/mask-length > [no-advertise]
[type < summary | nssa >]`

Use this command on a routing switch intended to operate as an ABR for the specified area to do either of the following:

- *Simultaneously create the area and corresponding range setting for routes to summarize or block.*
- *For an existing area, specify a range setting for routes to summarize or block.*

< ospf-area-id >: *Same area ID as on page 5-67 except you cannot use a backbone area number (0 or 0.0.0.0) for a stub area or NSSA.*

range < ip-addr/mask-length >: *Defines the range of route advertisements to either summarize for injection into the backbone area or to prevent from being injected into the backbone area.*

*The **ip-addr** value specifies the IP address portion of the range, and **mask-length** specifies the leftmost significant bits in the address. The ABR for the specified area compares the IP address of each outbound route advertisement with the address and significant bits in the mask to determine which routes to select for either summarizing or blocking. For example, a range of 10.10.32.1/14 specifies all routes in the range of 10.10.32.1 - 10.10.35.254.*

[no-advertise]: *Use this keyword only if you want to configure the ABR to prevent advertisement to the backbone of a specified range of routes. (This has the effect of “hiding” the specified range from the backbone area.) If you do not use this option, the ABR advertises the specified range of routes according to the **type < summary | nssa >** selection described below.*

[type < summary | nssa >]: *Configures the type of route summaries to advertise or block. If **type** is not used in the command, then the ABR defaults this setting to **summary**.*

summary: *Specifies internal routes in the configured range of route advertisements. If **no-advertise** (above) is used in the command, then the ABR prevents the selected internal routes from being summarized in a type-3 LSA and advertised to the backbone. If **no-advertise** is not used in the command, then the selected routes are summarized to the backbone in a type-3 LSA.*

nssa: *Specifies external routes (type-7 LSAs) in the configured range of route advertisements. If **no-advertise** (above) is used in the command, then the ABR prevents the selected external routes from being summarized in a type-5 LSA and advertised to the backbone. (Configure this option where an ABR for an NSSA advertises external routes that you do not want propagated to the backbone.) If **no-advertise** is not used in the command, then the selected routes learned from type-7 LSAs in the area are summarized to the backbone in a type-5 LSA.*

Examples of an ABR Allowing or Blocking Advertisement of a Range of Internal Routes Available in an Area. Both of the following commands define the same range of internal routes in area 30 to summarize for injection into the backbone area. (In this example, area 30 can be a normal or stub area, or an NSSA.)

```
ProCurve(ospf)# area 30 range 10.0.0.0/8
ProCurve(ospf)# area 30 range 10.0.0.0/8 type summary
```

Figure 5-28. Example of Defining a Range of Internal Routes To Advertise to the Backbone

For the same range of routes, you can use either of the following commands to block injection of a range of summary routes (type-3 LSAs) from area 30 into the backbone.

```
ProCurve(config)# area 30 range 10.0.0.0/8 type no-advertise
ProCurve(config)# area 30 range 10.0.0.0/8 type no-advertise summary
```

Figure 5-29. Example of Defining a Range of Internal Routes To Block from Advertising to the Backbone

Example of Allowing or Blocking a Range of External Routes Available Through an ASBR in an NSSA. This example applies only to external routes that can be advertised from an NSSA to the backbone.




Figure 5-30. Example of Allowing or Blocking a Range of External Route Advertisements to the Backbone

9. Optional: Influence Route Choices by Changing the Administrative Distance Default

The administrative distance value can be left in its default configuration setting unless a change is needed to improve OSPF performance for a specific network configuration.

The switch can learn about networks from various protocols, including RIP, and OSPF. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. For the switches, covered in this guide the administrative distance for OSPF routes is set at 110 for all route types (external, inter-area, and intra-area).

The switch selects one route over another based on the source of the route information. To do so, the switch can use the administrative distances assigned to the sources to influence route choices. You can change the distance settings in the OSPF global context to enable preference of one route type over another.

Syntax: distance < external | inter-area | intra-area > < 1 - 255 >

*Used in the OSPF configuration context to globally reconfigure the administrative distance priority for the specified route type. **1** is the highest priority; **255** is the lowest priority.*

external < 1 - 255 >: Changes the administrative distance for routes between the OSPF domain and other EGP domains.

inter-area < 1 - 255 >: Changes the administrative distance for routes between areas within the same OSPF domain.

intra-area < 1 - 255 >: Changes the administrative distance for routes within OSPF areas.

Default: 110; Range: 1 - 255

10: Optional: Change OSPF Trap Generation Choices

OSPF traps (defined by RFC 1850) are supported on the routing switches covered by this guide. OSPF trap generation is disabled by default, but you can use the following command to enable generation of any or all of the supported OSPF traps.

Syntax: [no] trap < trap-name | all >

*Used in the OSPF configuration context to enable or disable OSPF traps. The **no** form disables the specified trap.*

Default: All OSPF traps disabled.

all: *Enables or disables all OSPF traps available on the routing switch.*

trap-name: *Specifies a trap from table 5-8 to enable or disable.*

Table 5-8 summarizes OSPF traps supported on the switches covered in this guide, and their associated MIB objects from RFC 1850:

Table 5-8. OSPF Traps and Associated MIB Objects

OSPF Trap Name	MIB Object
interface-authentication-failure	ospflfAuthFailure
interface-config-error	ospflfConfigError
interface-receive-bad-packet	ospflfrxBadPacket
interface-retransmit-packet	ospfTxRetransmit
interface-state-change	
neighbor-state-change	ospfNbrStateChange
originate-lsa	ospfOriginateLsa
originate-maxage-lsa	ospfMaxAgeLsa
virtual-interface-authentication-failure	
virtual-interface-config-error	ospfVirtIfConfigError
virtual-interface-state-change	ospfVirtIfStateChange
virtual-neighbor-state-change	ospfVirtNbrStateChange
virtual-interface-receive-bad-packet	ospfVirtIfRxBad Packet
virtual-interface-retransmit-packet	ospfVirtIfTxRetransmit

For example, if you wanted to monitor the neighbor-state-change and interface-receive-bad-packet traps, you would use the following commands to configure the routing switch to enable the desired trap. The **show** command verifies the resulting OSPF trap configuration.

```
ProCurve(ospf)# trap neighbor-state-change
ProCurve(ospf)# trap interface-receive-bad-packet
ProCurve(ospf)# show ip ospf traps

OSPF Traps Enabled
=====

Neighbor State Change
Interface Receive Bad Packet
```

Figure 5-31. Example of Enabling OSPF Traps

11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings

The following OSPF interface parameters are automatically set to their default values. No change to the defaults is usually required unless needed for specific network configurations.

Parameter	Default	Page
cost	1	
dead-interval	40 seconds	
hello-interval	10 seconds	
priority	1	
retransmit-interval	5 seconds	
transit-delay	1 second	

A setting described in this section can be configured with the same value across all subnets in a VLAN or be configured on a per-interface basis with different values.

Note

Most of the parameters in this section also apply to virtual link configurations. However, when used on a virtual link configuration, the OSPF context requirement is different and the parameters are applied only to the interfaces included in the virtual link. Refer to “Optional: Adjust Virtual Link Performance by Changing the Interface Settings” on page 5-92.

Cost Per-Interface.

Syntax: ip ospf [ip-address | all] cost < 1 - 65535 >

Used in the VLAN context to indicate 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. Allows different costs for different subnets in the VLAN.

ip ospf cost < 1 - 65535 >: Assigns the specified cost to all networks configured on the VLAN.

ip ospf < ip-address > cost < 1 - 65535 >: Assigns the specified cost to the specified subnet on the VLAN.

ip ospf all cost < 1 - 65535 >: Assigns the specified cost to all networks configured on the VLAN. (Operates the same as the **ip ospf cost** option, above.)

Default: 1; Range 1 - 65535

Dead Interval Per-Interface.

Syntax: ip ospf [ip-address | all] dead-interval < 1 - 65535 >

Used in the VLAN context to indicate the number of seconds that a neighbor router waits for a hello packet from the specified interface before declaring the interface “down”. Allows different settings for different subnet interfaces in the VLAN.

ip ospf dead-interval < 1 - 65535 >: Assigns the specified dead interval to all networks configured on the VLAN.

ip ospf < ip-address > dead-interval < 1 - 65535 >: Assigns the specified dead interval to the specified subnet on the VLAN.

ip ospf all dead-interval < 1 - 65535 >: Assigns the specified dead interval to all networks configured on the VLAN. (Operates the same as the **ip ospf dead-interval** option, above.)

Default: 40 seconds; Range: 1 - 65535 seconds.

Hello Interval Per Interface.

Syntax: ip ospf [*ip-address* | all] hello-interval < 1 - 65535 >

Used in the VLAN context to indicate the length of time between the transmission of hello packets from the routing switch to adjacent neighbors. The value can be from 1 – 65535 seconds. The default is 10 seconds.. Allows different settings for different subnet interfaces in the VLAN.

ip ospf hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to all networks configured on the VLAN.*

ip ospf < ip-address > hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to the specified subnet on the VLAN.*

ip ospf all hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to all networks configured on the VLAN. (Operates the same as the ip ospf hello-interval option, above.)*

Default: 10 seconds; Range: 1 - 65535 seconds.

Priority Per-Interface.

Syntax: ip ospf [*ip-address* | all] priority < 1 - 255 >

Used in the VLAN context to enable changing the priority of an OSPF router. The priority is used when selecting the designated router (DR) and backup designated routers (BDRs). The value can be from 0 – 255 (with 255 as the highest priority). The default is 1. If you set the priority to 0, the routing switch does not participate in DR and BDR election. Allows different settings for different subnet interfaces in the VLAN.

ip ospf priority < 1 - 255 >: *Assigns the specified priority to all networks configured on the VLAN.*

ip ospf < ip-address > priority < 1 - 255 >: *Assigns the specified priority to the specified subnet on the VLAN.*

ip ospf all priority < 1 - 255 >: *Assigns the specified priority to all networks configured on the VLAN. (Operates the same as the ip ospf priority option, above.)*

Default: 1; Range: 0 - 255

Retransmit Interval Per-Interface.

Syntax: ip ospf [*ip-address* | all] retransmit-interval < 0 - 3600 >

Used in the VLAN context to enable changing the retransmission interval for link-state advertisements (LSAs) on an interface. The default is 5 seconds. Allows different settings for different subnet interfaces in the VLAN.

ip ospf priority < 1 - 255 >: Assigns the specified retransmit interval to all networks configured on the VLAN.

ip ospf < ip-address > priority < 1 - 255 >: Assigns the specified retransmit interval to the specified subnet on the VLAN.

ip ospf all priority < 1 - 255 >: Assigns the specified retransmit interval to all networks configured on the VLAN. (Operates the same as the **ip ospf priority** option, above.)

Default: 5 seconds; Range: 1 - 3600 seconds

Transit-Delay Per-Interface.

Syntax: ip ospf [*ip-address* | all] transit-delay < 1 - 3600 >

Used in the VLAN context to enable changing the time it takes to transmit Link State Update packets on this interface. Allows different settings for different subnet interfaces in the VLAN.

ip ospf transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the VLAN.

ip ospf < ip-address > transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the specified subnet on the VLAN.

ip ospf all transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the VLAN. (Operates the same as the **ip ospf transit-delay** option, above.)

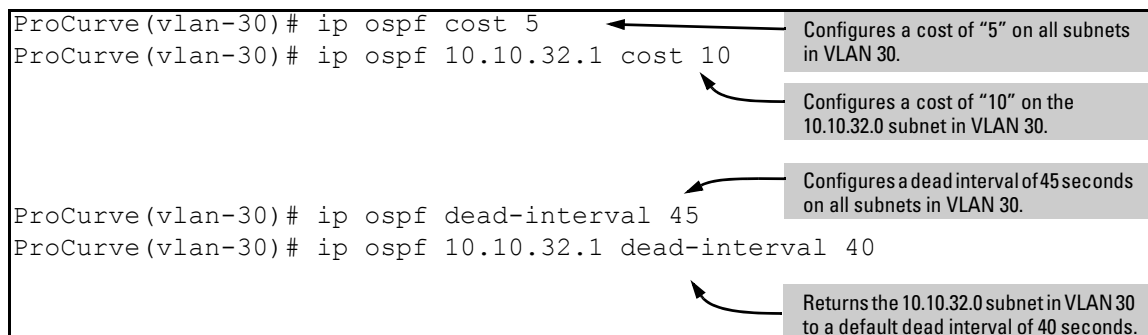
Default: 1 second; Range: 1 - 3600 seconds

Examples of Changing Per-Interface Settings. Suppose that VLAN 30 is multinetted, with two subnets in area 1 and one subnet in area 5:

```
vlan 30
 ip ospf 10.10.30.1 area 0.0.0.1
 ip ospf 10.10.31.1 area 0.0.0.1
 ip ospf 10.10.32.1 area 0.0.0.5
```

If you wanted to quickly reconfigure per-interface OSPF settings for VLAN 30, such as those listed below, you could use the commands shown in Figure 5-32.

- Assign a cost of “5” to the two subnets in area 1 and a cost of “10” to the subnet in area 5.
- Assign a dead interval of 45 seconds to the subnets in area 1 and retain the default setting (40 seconds) for the subnet in area 5.



```
ProCurve(vlan-30)# ip ospf cost 5
ProCurve(vlan-30)# ip ospf 10.10.32.1 cost 10

ProCurve(vlan-30)# ip ospf dead-interval 45
ProCurve(vlan-30)# ip ospf 10.10.32.1 dead-interval 40
```

Configures a cost of “5” on all subnets in VLAN 30.

Configures a cost of “10” on the 10.10.32.0 subnet in VLAN 30.

Configures a dead interval of 45 seconds on all subnets in VLAN 30.

Returns the 10.10.32.0 subnet in VLAN 30 to a default dead interval of 40 seconds.

Figure 5-32. Example of Reconfiguring Per-Interface Settings in a Multinetted VLAN

12. Optional: Configuring OSPF Interface Authentication

OSPF supports two methods of authentication for each VLAN or subnet—simple password and MD5. In addition, the value can be disabled, meaning no authentication is performed. Only one method of authentication can be active on a VLAN or subnet at a time, and if one method is configured on an interface, then configuring the alternative method on the same interface automatically overwrites the first method used. In the default configuration, OSPF authentication is disabled. All interfaces in the same network or subnet must have the same authentication method (password or MD5 key chain) and credentials.

OSPF Password Authentication.

Syntax: ip ospf [*ip-address*] authentication-key < *octet-string* >
no ip ospf [*ip-address*] authentication

*Used in the VLAN interface context to configure password authentication for all interfaces in the VLAN or for a specific subnet. The password takes effect immediately, and all OSPF packets transmitted on the interface contain this password. All OSPF packets received on the interface are checked for the password. If it is not present, then the packet is dropped. To disable password authentication on an interface, use the **no** form of the command.*

[*ip-address*]: *Used in subnetted VLAN contexts where you want to assign or remove a password associated with a specific subnet. Omit this option when you want the command to apply to all interfaces configured in the VLAN.*

< *octet-string* >: *An alphanumeric string of one to eight characters. (Spaces are not allowed.) To change the password, re-execute the command with the new password.*

Use **show ip ospf interface < *ip-address* >** to view the current authentication setting. (Refer to pages 5-103 and 5-105.)

Note: *To replace the password method with the MD5 method on a given interface, overwrite the password configuration by using the MD5 form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF password.)*

Default: Disabled

OSPF MD5 Authentication.

Syntax: ip ospf md5-auth-key-chain < chain-name-string >
no ip ospf [ip-address] authentication

*Used in the VLAN interface context to configure MD5 authentication for all interfaces in the VLAN or for a specific subnet. The MD5 authentication takes effect immediately, and all OSPF packets transmitted on the interface contain the designated key. All OSPF packets received on the interface are checked for the key. If it is not present, then the packet is dropped. To disable MD5 authentication on an interface, use the **no** form of the command.*

Note: Before using this authentication option, you must configure one or more key chains on the routing switch by using the Key Management System (KMS) described in the chapter titled “Key Management System” in the [Access Security Guide](#) for your routing switch

[ip-address]: Used in subnetted VLAN contexts where you want to assign or remove MD5 authentication associated with a specific subnet. Omit this option when you want the command to apply to all interfaces configured in the VLAN.

< chain-name-string >: The name of a key generated using the **key-chain < chain_name > key < key_id >** command. To change the MD5 authentication configured on an interface, re-execute the command with the new MD5 key.

Use **show ip ospf interface < ip-address >** to view the current authentication setting. (Refer to pages 5-103 and 5-105.)

Note: To replace the MD5 method with the password method on a given interface, overwrite the MD5 configuration by using the password form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF MD5 authentication.)

Default: Disabled

13. Configuring an ABR To Use a Virtual Link to the Backbone

All ABRs (area border routers) must have either a direct, physical or indirect, virtual link to the OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to the area backbone, the ABR can use a *virtual link* to provide a logical connection to another ABR having a direct physical connection to the area backbone. Both ABRs must belong to the same area, and this area becomes a transit area for traffic to and from the indirectly connected ABR.

Note

A backbone area can be purely virtual with no physical backbone links. Also note that virtual links can be “daisy chained”. If so, it may not have one end physically connected to the backbone.

Because both ABRs in a virtual link connection are in the same OSPF area, they use the same *transit area ID*. This setting is automatically determined by the ABRs and should match the area ID value configured on both ABRs in the virtual link.

The ABRs in a virtual link connection also identify each other with a *neighbor router* setting:

- On the ABR having the direct connection to the backbone area, the *neighbor router* is the IP address of the router interface needing a logical connection to the backbone.
- On the opposite ABR (the one needing a logical connection to the backbone), the *neighbor router* is the IP address of the ABR that is directly connected to the backbone.

Notes

By default, the router ID is the lowest numbered IP address or (user-configured) loopback interface configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 5-15.

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

Configuring a Virtual Link

Syntax: ip ospf area < area-id > virtual-link < ip-address >

Used on a pair of ABRs at opposite ends of a virtual link in the same area to configure the virtual link connection.

< area-id >: *This must be the same for both ABRs in the link, and is the area number of the virtual link transit area in either decimal or dotted decimal format.*

< ip-address >: *On an ABR directly connected to the backbone area, this value must be the IP address of an ABR (in the same area) needing a virtual link to the backbone area as a substitute for a direct physical connection. On the ABR that needs the virtual link to the backbone area, this value must be the IP address of the ABR (in the same area) having a direct physical connection to the backbone area.*

Example. Figure 5-33 shows an OSPF ABR, routing switch “A”, that lacks a direct connection to the backbone area (area 0). To provide backbone access to routing switch “A”, you can add a virtual link between routing switch “A” and routing switch “C”, using area 1 as a transit area. To configure the virtual link, define it on the routers that are at each end of the link. No configuration

for the virtual link is required on the other routers on the path through the transit area (such as routing switch “B” in this example).

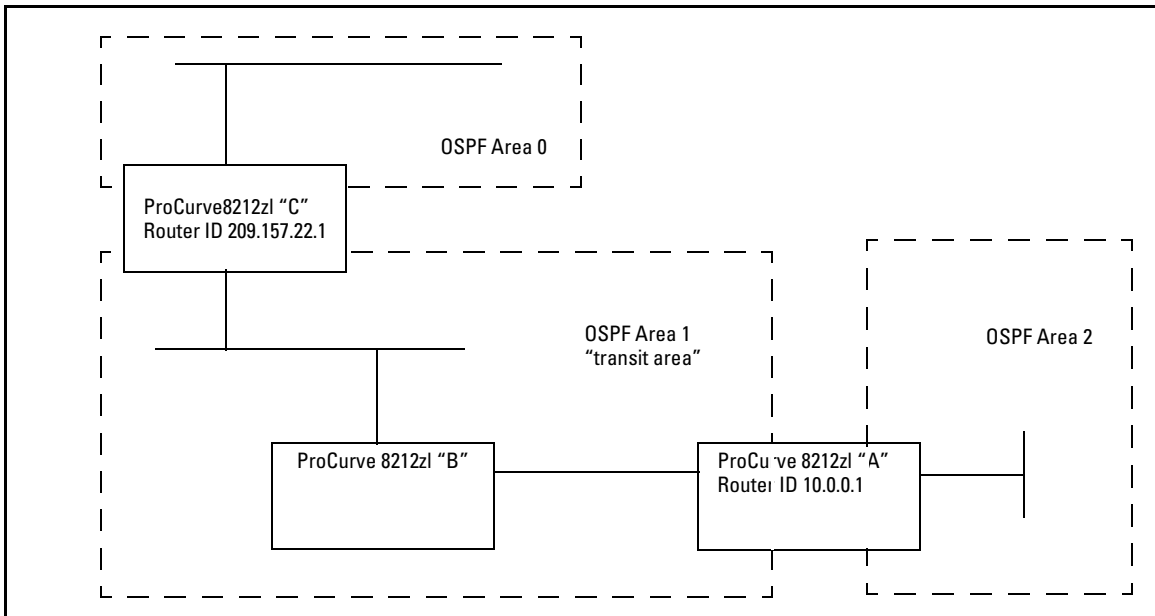


Figure 5-33. Defining OSPF virtual links within a network

To configure the virtual link on routing switch “A”, enter the following command specifying the area 1 interface on routing switch “C”:

```
ProCurve(ospf) # area 1 virtual-link 209.157.22.1
```

To configure the virtual link on routing switch “C”, enter the following command specifying the area 1 interface on routing switch “A”:

```
ProCurve(ospf) # area 1 virtual-link 10.0.0.1
```

Refer to “Optional: Adjust Virtual Link Performance by Changing the Interface Settings” on page 5-92 below for descriptions of virtual link interface parameters you can either use in their default settings or reconfigure as needed.

Optional: Adjust Virtual Link Performance by Changing the Interface Settings

The following OSPF interface parameters are automatically set to their default values for virtual links. No change to the defaults is usually required unless needed for specific network conditions. This is a subset of the parameters described under “11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings” on page 5-82. (The **cost** and **priority** settings are not configurable for a virtual link, and the commands for reconfiguring the settings are accessed in the router OSPF context instead of the VLAN context.)

Note

The parameter settings described in this section for virtual links must be the same on the ABRs at both ends of a given link.

Parameter	Default	Page
dead-interval	40 seconds	below
hello-interval	10 seconds	5-93
retransmit-interval	5 seconds	5-93
transit-delay	1 second	5-94

Dead Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > dead-interval < 1 - 65535 >

*Used in the router OSPF context on both ABRs in a virtual link to change the number of seconds that a neighbor router waits for a hello packet from the specified interface before declaring the interface “down”. This should be some multiple of the Hello interval. The **dead-interval** setting must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-90.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-114.)*

Default: 40 seconds; Range: 1 - 65535 seconds.

Hello Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > hello-interval < 1 - 65535 >

*Used in the router OSPF context on both ABRs in a virtual link to indicate the length of time between the transmission of hello packets between the ABRs on opposite ends of the virtual link. The value can be from 1 – 65535 seconds. The default is 10 seconds. The **hello-interval** setting must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-90.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-114.)*

Default: 10 seconds; Range: 1 - 65535 seconds.

Retransmit Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > retransmit-interval < 1 - 3600 >

*Used in the router OSPF context on both ABRs in a virtual link to change the number of seconds between link-state advertisement (LSA) retransmissions on the virtual link. The default is 5 seconds. The **retransmit-interval** setting must be the same on both ABRs on a given virtual link. This value is also used when retransmitting database description and link-state request packets.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-90.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-114.)*

Default: 5 seconds; Range: 1 - 3600 seconds

Transit-Delay on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > transit-delay < 0 - 3600 >

*Used in the router OSPF context on both ABRs in a virtual link to change the estimated number of seconds it takes to transmit a link state update packet over a virtual link. The **transit-delay** setting must be the same on both ABRs on a given virtual link.*

< area-id >: Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.

< ip-address >: For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of < ip-address > in the syntax description under “Configuring a Virtual Link” on page 5-90.)

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-114.)*

Default: 1 second; Range: 1 - 3600 seconds

Example. To change the hello-interval on the virtual link configured for the network in figure 5-33 on page 5-91 to 60 seconds:

- On routing switch “A” (IP address 10.0.0.1) you would use the following command to reconfigure the current hello-interval to 60 seconds:

```
ProCurve(ospf)# area 1 virtual-link 209.157.22.1  
hello-interval 60
```

- On routing switch “C” (IP address 209.157.22.1) you would use the following command to reconfigure the current hello-interval to 60 seconds

```
ProCurve(ospf)# area 1 virtual-link 10.0.0.1  
hello-interval 60
```

Configuring OSPF Authentication on a Virtual Link

OSPF supports the same two methods of authentication for virtual links as it does for VLANs and subnets in an area—password and MD5. In the default configuration, OSPF authentication is disabled. Only one method of authentication can be active on a virtual link at a time, and if one method is configured on a virtual link, then configuring the alternative method on the same link automatically replaces the first method with the second. Both ends of a virtual link must use the same authentication method (none, password, or MD5 key chain) and related credentials. (Any interfaces that share a VLAN or subnet with the interface used on an ABR for a virtual link, including intermediate routing switches, must be configured with the same OSPF authentication.)

OSPF Password Authentication on a Virtual Link.

Syntax: `area < area-id > virtual-link < ip-addr > authentication-key < octet-string >`
`no area 1 virtual-link < ip-address > authentication`

*Used to configure password authentication in the router OSPF context on both ABRs in a virtual link. The password takes effect immediately, and all OSPF packets transmitted on the link contain this password. Every OSPF packet received on the interface for the virtual link on each ABR is checked for the password. If it is not present, then the packet is dropped. To disable password authentication on an ABR interface used for a virtual link, use the **no** form of the command. The password must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-addr >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-90.)*

< octet-string >: *An alphanumeric string of one to eight characters. (Spaces are not allowed.) To change the password, re-execute the command with the new password.*

Note: *To replace the password method with the MD5 method on a given interface, overwrite the password configuration by using the MD5 form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF password.)*

Default: Disabled

OSPF MD5 Authentication on a Virtual Link.

Syntax: ip ospf md5-auth-key-chain < chain-name-string >
no ip ospf [ip-address] authentication

*Used to configure MD5 authentication in the router OSPF context on both ABRs in a virtual link . The MD5 authentication takes effect immediately, and all OSPF packets transmitted on the link contain the designated key. Every OSPF packet received on the interface for the virtual link on each ABR is checked for the key. If it is not present, then the packet is dropped. To disable MD5 authentication on an ABR interface used for a virtual link, use the **no** form of the command. The password must be the same on both ABRs on a given virtual link.*

Note: *Before using this authentication option, you must configure one or more key chains on the routing switch by using the Key Management System (KMS) described in the chapter titled “Key Management System” in the [Access Security Guide](#) for your routing switch.*

[ip-address]: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of < ip-address > in the syntax description under “Configuring a Virtual Link” on page 5-90.)*

< chain-name-string >: *The name of a key generated using the **key-chain < chain_name > key < key_id >** command. To change the MD5 authentication configured on a virtual link, re-execute the command with the new MD5 key.*

Note: *To replace the MD5 method with the password method on a virtual link, overwrite the MD5 configuration by using the password form of the command shown under “OSPF Password Authentication on a Virtual Link” on page 5-95. (It is not necessary to disable the currently configured OSPF MD5 authentication.)*

Default: Disabled

OSPF Passive

OSPF sends link-state advertisements (LSAs) to all other routers in the same Autonomous System (AS). To limit the flooding of LSAs throughout the AS you can configure OSPF to be passive. OSPF does not run in the AS, but it does advertise the interface as a stub link into OSPF. Routing updates are accepted by a passive interface, but not sent out.

There is a limit of 512 total active and passive interfaces, but only a total of 128 can be active interfaces.

To configure a passive OSPF interface, enter this command in vlan context:

```
ProCurve(vlan-1)# ip ospf passive
```

Syntax: [no] ip ospf <ip-addr> passive

Configures passive OSPF for an Autonomous System.

*The **no** option disables the passive option; the interface becomes an active interface.*

Default: Active

<ip-addr>: *Optionally you can configure an IP address on the VLAN*

To display the OSPF passive information, enter the command shown in Figure 5-34:

```
ProCurve(vlan-1)# show ip ospf interface
```

OSPF Interface Status							
IP Address	Status	Area ID	State	Auth-type	Cost	Priority	Passive
-----	-----	-----	-----	-----	-----	-----	-----
10.10.10.1	enabled	0.0.0.2	down	none	1	1	Yes
10.12.13.1	enabled	0.0.0.2	wait	none	1	1	No

Figure 5-34. Example of the show ip ospf interface Command with Passive Configured on an Interface

You can display the OSPF passive information for a particular VLAN, as shown in Figure 5-35.

```
ProCurve(config) show ip ospf interface vlan 4
OSPF configuration and statistics for VLAN 4
OSPF Interface Status for 10.10.10.1
IP Address:      : 10.10.10.1   Status  : enabled
AreaID          : 0.0.0.2     Passive : Yes
State           : DOWN
Cost            : 1
Type            : BCAST
Auth-type       : none
Chain           :
Priority        : 1
Transit Delay   : 1
Hello Interval  : 10
Designated Router:
Backup Desig. Rtr:
Retrans Interval : 5
Rtr Dead Interval : 40
Events           : 0
Passive          : yes
```

Figure 5-35. Example of the show ip ospf interface Command for a specific VLAN with Passive Configured on an Interface

Displaying OSPF Information

You can use CLI commands to display the following OSPF information:

OSPF Information Type	Page
General Information	5-99
Area information	5-100
External link state information	5-101
Interface information	5-103
Link state information	5-106
Neighbor information	5-109
Route information	5-115
Virtual Neighbor information	5-112
Virtual Link information	5-113
OSPF Traps enabled	5-117

Displaying General OSPF Configuration Information

To display general OSPF configuration information, enter **show ip ospf general** at any CLI level:

```

ProCurve# show ip ospf general

OSPF General Status

  OSPF protocol                : enabled
  Router ID                    : 10.0.8.36
  RFC 1583 compatibility       : compatible

  Intra-area distance          : 110
  Inter-area distance          : 110
  AS-external distance         : 110

  Default import metric        : 1
  Default import metric type   : external type 2

  Area Border                  : yes
  AS Border                    : yes
  External LSA Count           : 9
  External LSA Checksum Sum    : 408218
  Originate New LSA Count      : 24814
  Receive New LSA Count        : 14889
  
```

Figure 5-36. Example of Show IP OSPF General Output

Syntax: show ip ospf general

The following fields are shown in the OSPF general status display:

Table 5-9. CLI Display of OSPF General Information

Field	Content
OSPF protocol	whether OSPF is currently enabled.
Router ID	the Router ID that this routing switch is currently using to identify itself
RFC 1583 compatibility	whether the routing switch is currently using RFC 1583 (compatible) or RFC 2328 (non-compatible rules for calculating external routes).
Intra-area distance	the administrative distance for routes within OSPF areas

Field	Content
Inter-area distance	the administrative distance for routes between areas within the same OSPF domain
AS-external	the administrative distance for routes between the OSPF domain and other, Exterior Gateway Protocol domains
Default import metric	the default metric that will be used for any routes redistributed into OSPF by this routing switch
Default import metric type	the metric type (type 1 or type 2) that will be used for any routes redistributed into OSPF by this routing switch
Area Border	whether this routing switch is currently acting as an area border router
AS Border	whether this routing switch is currently acting as an autonomous system border router (redistributing routes)
External LSA Count	the total number of external LSAs currently in the routing switch's link state database
External LSA Checksum Sum	the sum of the checksums of all external LSAs currently in the routing switch's link state database (quick check for whether database is in sync with other routers in the routing domain)
Originate New LSA Count	count of the number of times this switch has originated a new LSA
Receive New LSA Count	count of the number of times this switch has received a new LSA

Displaying OSPF Area Information

To display OSPF area information, enter **show ip ospf area** at any CLI level:

```
ProCurve(config)# show ip ospf area

OSPF Area Information

Area ID          Type    Cost  SPFR  ABR   ASBR  LSA   Checksum
-----
0.0.0.0          normal  0      1      0     0     1     0x0000781f
192.147.60.0     normal  0      1      0     0     1     0x0000fee6
192.147.80.0     stub    1      1      0     0     2     0x000181cd
```

Figure 5-37. Example of Show IP OSPF Area Output

Syntax: show ip ospf area [*ospf-area-id*]

The `[ospf-area-id]` parameter shows information for the specified area. If no area is specified, information for all the OSPF areas configured is displayed.

The OSPF area display shows the following information:

Table 5-10. CLI Display of OSPF Area Information

Field	Content
Area ID	The identifier for this area.
Type	The area type, which can be either “normal” or “stub”.
Cost	The metric for the default route that the routing switch will inject into a stub area if the routing switch is an ABR for the area. This value only applies to stub areas.
SPFR	The number of times the routing switch has run the shortest path first route calculation for this area.
ABR	The number of area border routers in this area.
ASBR	The number of autonomous system border routers in this area.
LSA	The number of LSAs in the link state database for this area.
Chksum(Hex)	The sum of the checksums of all LSAs currently in the area’s link state database. This value can be compared to the value for other routers in the area to verify database synchronization.

Displaying OSPF External Link State Information

To display external link state information, enter **show ip ospf external-link-state** at any CLI level. When you enter this command, an output similar to the following is displayed:

```
ProCurve# show ip ospf external-link-state

OSPF External LSAs

Link State ID   Router ID      Age   Sequence #   Checksum
-----
10.3.7.0        10.0.8.37     232  0x80000005  0x0000d99f
10.3.8.0        10.0.8.37     232  0x80000005  0x0000cea9
10.3.9.0        10.0.8.37     232  0x80000005  0x0000c3b3
10.3.10.0       10.0.8.37     232  0x80000005  0x0000b8bd
10.3.33.0       10.0.8.36     1098 0x800009cd  0x0000b9dd
```

Figure 5-38. Example of Show IP OSPF External-Link-State Output

Syntax: show ip ospf external-link-state

The OSPF external link state display shows the following information:

Table 5-11. CLI Display of OSPF External Link State Information

Field	Content
Link State ID	LSA ID for this LSA. Normally, the destination of the external route, but may have some "host" bits set.
Router ID	Router ID of the router that originated this external LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Syntax show ip ospf external-link-state [status] [*subset-options*]
:

router-id < *ip-addr* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified router ID only. Can also be filtered by using the **link-state-id** or **sequence-number** options.*

sequence-number < *integer* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified sequence number. Can also be filtered by using the **link-state-id** or **router-id** options.*

link-state-id < *ip-addr* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified ID only. Can also be filtered by using the **sequence-number** or **router-id** options.*

Syntax: show ip ospf external-link-state [status] advertise

*Displays the hexadecimal data in the specified LSA packet, the actual contents of the LSAs. Can also be filtered by using the **link-state-id**, **router-id**, or **sequence-number** options.*

```

ProCurve# show ip ospf external-link-state advertise

OSPF External LSAs
Advertisements
-----
000302050a0307000a00082580000005d99f0024ffffff008000000a0000000000000000
000302050a0308000a00082580000005cea90024ffffff008000000a0000000000000000
000302050a0309000a00082580000005c3b30024ffffff008000000a0000000000000000
000302050a030a000a00082580000005b8bd0024ffffff008000000a0000000000000000
000002050a0321000a000824800009cdb9dd0024ffffff00800000010000000000000000

```

Figure 5-39. Example of the Output for Show IP OSPF External-Link-State Advertise

Displaying OSPF Interface Information

To display OSPF interface information, enter **show ip ospf interface** at any CLI level:

```

ProCurve# show ip ospf interface

OSPF Interface Status

IP Address      Status   Area ID      State   Auth-type  Cost   Priority
-----
10.3.18.36     enabled  10.3.16.0    BDR    none       1      1
10.3.53.36     enabled  10.3.48.0    BDR    none       1      1

```

Figure 5-40. Example of the Output for Show IP OSPF Interface

Syntax: show ip ospf interface [vlan <vlan-id> | <ip-addr>]

The OSPF interface display shows the following information:

Table 5-12. CLI Display of OSPF Interface Information

Field	Content
IP Address	The local IP address for this interface.
Status	enabled or disabled - whether OSPF is currently enabled on this interface.
Area ID	The ID of the area that this interface is in.

Field	Content
State	The current state of the interface. The value will be one of the following: <ul style="list-style-type: none">• DOWN - the underlying VLAN is down• WAIT - the underlying VLAN is up, but we are waiting to hear hellos from other routers on this interface before we run designated router election• DR - this switch is the designated router for this interface• BDR - this switch is the backup designated router for this interface• DROTHER - this router is not the designated router or backup designated router for this interface
Auth-type	none or simple - will be none if no authentication key is configured, simple if an authentication key is configured. All routers running OSPF on the same link must be using the same authentication type and key.
Chain	The name of the key chain configured for the specified interface. (Refer to the chapter titled "Key Management System" in the <i>Access Security Guide</i> for your routing switch.
Cost	The OSPF's metric for this interface.
Priority	This routing switch's priority on this interface for use in the designated router election algorithm.

The **< ip-addr >** parameter displays the OSPF interface information for the specified IP address.

The **< vlan-id >** parameter displays the OSPF interface information for the specified IP address.

Displaying OSPF Interface Information for a Specific VLAN or IP Address

To display OSPF interface information for a specific VLAN or IP address, enter **show ip ospf interface < ip-addr >** at any CLI level. For example:

```
ProCurve(ospf)# sho ip ospf int 10.10.50.1

OSPF Interface Status for 10.3.1836

IP Address       : 10.3.18.36           Status  : enabled
Area ID         : 10.3.16.0

State  : BDR                          Auth-type : none
Cost   : 1                             Chain     :
Type   : BCAST                         Priority  : 1

Transit Delay   : 1                    Retrans Interval : 5
Hello Interval  : 10                    Rtr Dead Interval : 40
Designated Router : 10.3.18.34          Events           : 3
Backup Desig. Rtr : 10.3.18.36
```

Figure 5-41. Example of Show IP OSPF Interface < ip-addr > Output

Syntax: show ip ospf interface [vlan < vlan-id > | < ip-addr >]

The OSPF interface display for a specific VLAN or IP address has the same information as the non-specific **show ip ospf interface** command for the **IP Address, Area ID, Status, State, Auth-type, Cost, and Priority** fields. See the information for the general command in table 5-12 on page 5-103 for definitions of these fields.

The **show ip ospf interface** command for a specific VLAN or IP address shows the following additional information:

Table 5-13. CLI Display of OSPF Interface Information – VLAN or IP Address

Field	Content
Type	Will always be BCAST for interfaces on this routing switch. Point-to-point or NBMA (frame relay or ATM) type interfaces are not supported on the switches covered in this guide.
Transit Delay	Configured transit delay for this interface.

Field	Content
Retrans Interval	Configured retransmit interval for this interface.
Hello Interval	Configured hello interval for this interface.
Rtr Dead Interval	Configured router dead interval for this interface.
Designated Router	IP address of the router that has been elected designated router on this interface.
Backup Desig. Rtr	IP address of the router that has been elected backup designated router on this interface.
Events	Number of times the interface state has changed.

If you use **show ip ospf interface vlan < vlan-id >**, the output will be the same as shown in the previous table, but for the IP address on the indicated VLAN.

Displaying OSPF Link State Information

To display OSPF link state information, enter **show ip ospf link-state** at any CLI level. When you enter this command, the switch displays an output similar to the following for all configured areas:

```

OSPF Link State Database for Area 0.0.0.0
      Advertising
LSA Type  Link State ID  Router ID      Age  Sequence #  Checksum
-----
Router    10.0.8.32             10.0.8.32      65   0x80000281  0x0000a7b6
Router    10.0.8.33             10.0.8.33     1638 0x80000005  0x0000a7c8
Network   10.3.2.37             10.0.8.37     1695 0x80000006  0x00000443
Summary   10.3.16.0             10.0.8.33     1638 0x80000007  0x0000c242
Summary   10.3.16.0             10.0.8.35     1316 0x80000008  0x0000aa58
Summary   10.3.17.0             10.0.8.33     1638 0x8000027b  0x0000becf
Summary   10.3.17.0             10.0.8.35     1316 0x80000008  0x0000a957
AsbSummary 10.0.8.36             10.0.8.33     1412 0x80000002  0x00002cba

OSPF Link State Database for Area 10.3.16.0
      Advertising
LSA Type  Link State ID  Router ID      Age  Sequence #  Checksum
-----
Router    10.0.8.33             10.0.8.33     1727 0x8000027e  0x0000d53c
Router    10.0.8.34             10.0.8.34     1420 0x80000283  0x0000de4f
Network   10.3.16.34          10.0.8.34     1735 0x80000005  0x00001465

```

Figure 5-42. Example of Show IP OSPF Link-State Output

The OSPF link state display shows the following contents of the LSA database; one table for each area:

Table 5-14. CLI Display of OSPF Link State Information

Field	Content
LSA Type	Type of LSA. The possible types are: <ul style="list-style-type: none"> • Router • Summary • Network • AsbSummary
Link State ID	LSA ID for this LSA. The meaning depends on the LSA type.
Advertised Router ID	Router ID of the router that originated this LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Syntax show ip ospf link-state [status] [subset-options] [< advertise [subset-options] >]

advertise: *Displays the hexadecimal data in LSA packets (advertisements) for the OSPF area(s) configured on the routing switch. The output can also be filtered by area (**area-id**), **link-state-id**, **router-id**, **sequence-number**, and/or **type**.*

Default: All OSPF areas configured on the routing switch.

ospf-area-id: *Used to restrict display of LSA database or advertisements to show only the data from a specific OSPF area. Can also be used with other subset options (**router-id**, **sequence-number**, **external link-state-id**, and/or **type**) to further define the source of displayed information.*

link-state-id < ip-addr >

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified IP address as a link-state ID. Can also be used with other subset options (**ospf-area-id**, **router-id**, **sequence-number**, **external link-state-id**, and **type**) to further define the source of displayed information.*

`router-id < ip-addr >`

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified router ID. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **sequence-number**, and **type**) to further define the source of displayed information.*

`sequence-number < integer >`

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified sequence number. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **router-id**, and **type**) to further define the source of displayed information.*

`type < router | network | summary | as-summary | external | multicast | nssa >`

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified type. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **router-id**, and **sequence-number**) to further define the source of displayed information.*

An example of **show ip ospf link-state advertise** output is:

```

ProCurve_8212(config)# show ip ospf link-state advertise

OSPF Link State Database for Area 0.0.0.0

Advertisements
-----
000202010a0008200a00082080000281a7b60054000000050a030e00ffffff0003000001...
000202010a0008210a00082180000006a5c90024010000010a0008230a03112104000002
000102010a0008230a00082380000015755d006c010000070a030600ffffff0003000001...
000202020a0302250a0008258000000702440024ffffff000a0008250a0008230a000820
000202030a0310000a00082180000008c043001cffffff0000000002
000102030a0310000a00082380000009a859001cffffff0000000001
000002030a0310000a00082480000009ac53001cffffff0000000002
000202040a0008240a000821800000032abb001c000000000000000b
000102040a0008240a00082380000004c12a001c0000000000000002

OSPF Link State Database for Area 10.3.16.0

Advertisements
-----
000202010a0008210a0008218000027fd33d00540500000050a031900ffffff0003000001...
000102010a0008220a00082280000284dc500060000000060a031500ffffff0003000001...
000102020a0311220a0008228000027bf9080020ffffff000a0008220a000821

```

Figure 5-43. Example of the Output for Show IP OSPF Link-State Advertise

Displaying OSPF Neighbor Information

To display OSPF information for all neighbors, enter **show ip ospf neighbor** at any CLI level:

```

OSPF Neighbor Information

Router ID      Pri  IP Address      NbIfState  State      Rxmt  QLen  Events
-----
10.0.8.34      1    10.3.18.34      DR          FULL       0      6      6
10.3.53.38     1    10.3.53.38      DR          FULL       0      6      6

```

Figure 5-44. Example of Show IP OSPF Neighbor Output

Syntax: show ip ospf neighbor [*ip-addr*]

The [*ip-addr*] can be specified to retrieve detailed information for the specific neighbor only. This is the IP address of the neighbor, not the router ID.

This display shows the following information.

Table 5-15. CLI Display of OSPF Neighbor Information

Field	Description
Router ID	The router ID of the neighbor.
Pri	The OSPF priority of the neighbor. The priority is used during election of the Designated Router (DR) and Backup designated Router (BDR).
IP Address	The IP address of this routing switch's interface with the neighbor.
NbIfState	The neighbor interface state. The possible values are: <ul style="list-style-type: none"> • DR – this neighbor is the elected designated router for the interface. • BDR – this neighbor is the elected backup designated router for the interface. • blank – this neighbor is neither the DR or the BDR for the interface.
State	The state of the conversation (the adjacency) between your routing switch and the neighbor. The possible values are: <ul style="list-style-type: none"> • INIT – A Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor. (The switch itself did not appear in the neighbor's Hello packet.) All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface. • 2WAY – Communication between the two routers is bidirectional. This is the most advanced state before beginning adjacency establishment. The Designated Router and Backup Designated Router are selected from the set of neighbors in the 2Way state or greater. • EXSTART – The first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description (DD) sequence number. Neighbor conversations in this state or greater are called adjacencies. • EXCHANGE – The switch is describing its entire link state database by sending Database Description packets to the neighbor. Each Database Description packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description packet can be outstanding at any time. In this state, Link State Request packets can also be sent asking for the neighbor's more recent advertisements. All adjacencies in Exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets. • LOADING – Link State Request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the Exchange state. • FULL – The neighboring routers are fully adjacent. These adjacencies will now appear in router links and network link advertisements.
Rxmt QLen	Remote transmit queue length – the number of LSAs that the routing switch has sent to this neighbor and for which the routing switch is awaiting acknowledgements.
Events	The number of times the neighbor's state has changed.

Displaying OSPF Redistribution Information

As described under “2. Enable Route Redistribution” on page 5-76, you can configure the routing switch to redistribute connected, static, and RIP routes into OSPF. When you redistribute a route into OSPF, the routing switch can use OSPF to advertise the route to its OSPF neighbors.

To display the status of the OSPF redistribution, enter **show ip ospf redistribute** at any CLI context level:

```
ProCurve# show ip ospf redistribute

OSPF redistributing
Route type Status
-----
connected  enabled
static      enabled
rip         enabled
```

Figure 5-45. Example of Output for Show IP OSPF Redistribute

The display shows whether redistribution of each of the route types, connected, static, and RIP is enabled.

Displaying OSPF Redistribution Filter (restrict) Information

As described under “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-74, you can configure the redistribution filters on the routing switch to restrict route redistribution by OSPF.

To display the status of the OSPF redistribution filters, enter **show ip ospf restrict** at any CLI context level.

```
ProCurve# show ip ospf restrict

OSPF restrict list

IP Address      Mask
-----
10.0.8.0        255.255.248.0
15.0.0.0        255.0.0.0
```

Figure 5-46. Example of Output for Show IP OSPF Restrict

This display shows the configured restrict entries.

Displaying OSPF Virtual Neighbor Information

If virtual links are configured on the routing switch, you can display OSPF virtual neighbor information by entering **show ip ospf virtual-neighbor** at any CLI level.

```
OSPF Virtual Interface Neighbor Information
```

Router ID	Area ID	State	IP Address	Events
10.0.8.33	10.3.16.0	FULL	10.3.17.33	5
10.0.8.36	10.3.16.0	FULL	10.3.18.36	5

Figure 5-47. Example of Output for Show IP OSPF Virtual-Neighbor

Syntax: show ip ospf virtual-neighbor [area < *area-id* > | < *ip-address* >]

This display shows the following information.

Table 5-16. CLI Display of OSPF Virtual Neighbor Information

Field	Description
Router ID	The router ID of this virtual neighbor (configured).
Area ID	The area ID of the transit area for the virtual link to this neighbor (configured).
State	The state of the adjacency with this virtual neighbor. The possible values are the same as the OSPF neighbor states. See the State parameter definition in table 5-15 on page 5-110. Note that virtual neighbors should never stay in the 2WAY state.
IP Address	IP address of the virtual neighbor that the routing switch is using to communicate to that virtual neighbor.
Events	The number of times the virtual neighbor's state has changed.

Notice from the syntax statement that **ip-address** can be specified to display detailed information for a particular virtual neighbor. If an **area-id** is specified only virtual neighbors belonging to that area are shown.

Displaying OSPF Virtual Link Information

If virtual links are configured on a routing switch, you can display OSPF virtual link information by entering **show ip ospf virtual-link** at any CLI level.

```
ProCurve# show ip ospf virtual-link

OSPF Virtual Interface Status

  Transit AreaID  Neighbor Router  Authentication  Interface State
  -----
  10.3.16.0       10.0.8.33       none            P2P
  10.3.16.0       10.0.8.36       none            P2P
```

Figure 5-48. Example of Output for Show IP OSPF Virtual-Link

Syntax: show ip ospf virtual-link [area < *area-id* > | < *ip-address* >]

This display shows the following information.

Table 5-17. CLI Display of OSPF Virtual Link Information

Field	Description
Transit Area ID	Area ID of transit area for the virtual link.
Neighbor Router	Router ID of the virtual neighbor.
Authentication	none or simple (same as for normal interface).
Interface State	The state of the virtual link to the virtual neighbor. The possible values are: <ul style="list-style-type: none"> DOWN – the routing switch has not yet found a route to the virtual neighbor. P2P – (point-to-point) the routing switch has found a route to the virtual neighbor. Virtual links are “virtual” serial links, hence the point-to-point terminology.

Notice from the syntax statement that ***ip-address*** can be specified to display detailed information for a particular virtual neighbor. If an ***area-id*** is specified only virtual links belonging to that area are shown.

Example:. To get OSPF virtual link information for IP address 10.0.8.33, enter **show ip ospf virtual-link 10.0.8.33**. A display similar to the following is shown.

```
ProCurve# show ip ospf virtual-link 10.0.8.33

OSPF Virtual Interface Status for interface 10.0.8.33
  Transit AreaID   : 10.3.16.0
  Neighbor Router  : 10.0.8.33

  Authentication   : none                Chain           :
  Interface State  : P2P                 Transit Delay    : 1
  Events           : 1                   Rtr Interval    : 5
  Dead Interval    : 40                   Hello Interval   : 10
```

Figure 5-49. Example of Output for Show IP OSPF Virtual-Link < ip-addr >

In this display, these fields show the same type of information as described for the general OSPF virtual link display: **Transit Area ID**, **Neighbor Router**, **Authentication**, and **Interface State**. This display shows the following additional information:

Table 5-18. CLI Display of OSPF Virtual Link Information – Specific IP Address

Field	Description
Events	The number of times the virtual link interface state has changed.
Transit delay	The configured transit delay for the virtual link.
Rtr Interval	The configured retransmit interval for the virtual link.
Hello Interval	The configured hello interval for the virtual link.
Dead Interval	The configured router dead interval for the virtual link

Displaying OSPF Route Information

To display OSPF route and other OSPF configuration information, enter **show ip ospf** at any CLI level:

```
ProCurve# show ip ospf

OSPF Configuration Information

  OSPF protocol   : enabled
  Router ID      : 10.0.8.35

Currently defined areas:
```

Area ID	Type	Stub Default Cost	Stub Summary LSA	Stub Metric Type
backbone	normal	1	don't send	ospf metric
10.3.16.0	normal	1	don't send	ospf metric
10.3.32.0	normal	1	don't send	ospf metric

```
Currently defined address ranges:
```

Area ID	LSA Type	IP Network	Network Mask	Advertise
10.3.16.0	Summary	10.3.16.0	255.255.255.0	yes

```
OSPF interface configuration:
```

IP Address	Area ID	Admin Status	Type	Authen Type	Cost	Pri
10.3.2.35	backbone	enabled	BCAST	none	1	1
10.3.3.35	backbone	enabled	BCAST	none	1	1
10.3.16.35	10.3.16.0	enabled	BCAST	none	1	1
10.3.32.35	10.3.32.0	enabled	BCAST	none	1	1

```
OSPF configured interface timers:
```

IP Address	Transit Delay	Retransmit Interval	Hello Interval	Dead Interval
10.3.2.35	1	5	10	40
10.3.3.35	1	5	10	40
10.3.16.35	1	5	10	40
10.3.32.35	1	5	10	40

```
OSPF configured virtual interfaces:
```

Area ID	Router ID	Authen Type	Xmit Delay	Rxmt Intvl	Hello Intvl	Dead Interval
10.3.16.0	10.0.8.33	none	1	5	10	40
10.3.16.0	10.0.8.36	none	1	5	10	40

Figure 5-50. Example of Output for Show IP OSPF

Syntax: show ip ospf

This screen has a lot of information, most of it already covered in other show commands. The following table shows definitions for the fields:

Table 5-19. CLI Display of OSPF Route and Status Information

Field	Description
OSPF protocol	enabled or disabled – indicates if OSPF is currently enabled.
Router ID	The Router ID that this routing switch is currently using to identify itself.
Currently Defined Areas:	
Area ID	The identifier for this area.
Type	The type of OSPF area (normal or stub).
Stub Default Cost	The metric for any default route we injected into a stub area if the routing switch is an ABR for the area. This value only applies to stub areas.
Stub Summary LSA	send or don't send – indicates the state of the no-summary option for the stub area. The value indicates if the area is “totally stubby” (no summaries sent from other areas) or just “stub” (summaries sent). Only applies to stub areas, and only takes effect if the routing switch is the ABR for the area.
Stub Metric Type	This value is always ospf metric .
Currently defined address ranges:	
Area ID	The area where the address range is configured.
LSA Type	This value is always Summary .
IP Network	The address part of the address range specification.
Network Mask	The mask part of the address range specification.
Advertise	Whether we are advertising (yes) or suppressing (no) this address range.

Note

The remaining interface and virtual link information is the same as for the previously described OSPF show commands. Refer to Table 5-12 (page 5-103) and Table 5-13 (page 5-105).

Displaying OSPF Traps Enabled

In the default configuration, OSPF traps are disabled. Use this command to view which OSPF traps have been enabled.

Syntax: show ip ospf traps

Lists the OSPF traps currently enabled on the routing switch.

For more information on OSPF trap use, refer to “10: Optional: Change OSPF Trap Generation Choices” on page 5-81.

OSPF Equal-Cost Multipath (ECMP) for Different Subnets Available Through the Same Next-Hop Routes

The switches covered by this guide support optional load-sharing across redundant links where the network offers two, three, or four equal-cost next-hop routes for traffic to different subnets. (All traffic for different hosts in the same subnet goes through the same next-hop router.)

For example, in the OSPF network shown below, IP load-sharing is enabled on router “A”. In this case, OSPF calculates three equal-cost next-hop routes for each of the subnets and then distributes per-subnet route assignments across these three routes.

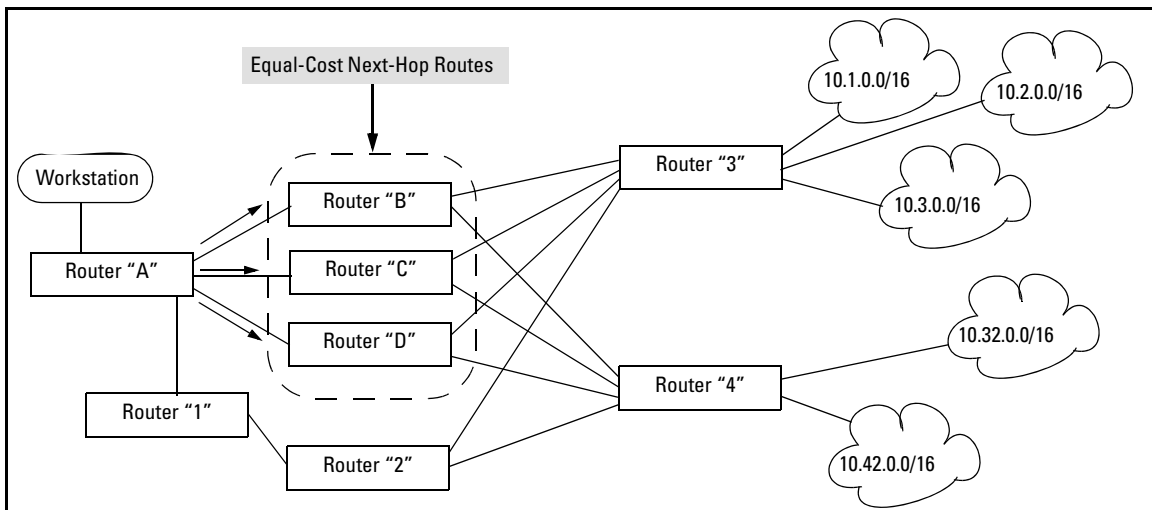


Figure 5-51. Example of Load-Sharing Traffic to Different Subnets Through Equal-Cost Next-Hop Routers

Example of a Routing Table for the Network in Figure 5-51

Destination Subnet	Router "A" Next Hop
10.1.0.0/16	Router "C"
10.2.0.0/16	Router "D"
10.3.0.0/16	Router "B"
10.32.0.0/16	Router "B"
10.42.0.0/16	Router "D"

Note that IP load-sharing does not affect routed traffic to different hosts on the same subnet. That is, all traffic for different hosts on the same subnet will go through the same next-hop router. For example, if subnet 10.32.0.0 includes two servers at 10.32.0.11 and 10.32.0.22, then all traffic from router "A" to these servers will go through router "B".

Syntax: *[no] ip load-sharing < 2 - 4 >*

*When OSPF is enabled and multiple, equal-cost, next-hop routes are available for traffic destinations on different subnets, this feature, by default, enables load-sharing among up to four next-hop routes. The **no** form of the command disables this load-sharing so that only one route in a group of multiple, equal-cost, next-hop routes is used for traffic that could otherwise be load-shared across multiple routes. For example, in figure 5-51 on page 5-117, the next-hop routers "B", "C", and "D" are available for equal-cost load-sharing of eligible traffic. Disabling IP load-sharing means that router "A" selects only one next-hop router for traffic that is actually eligible for load-sharing through different next-hop routers. (Default: Enabled with four equal-cost, next-hop routes allowed)*

Note: *In the default configuration, IP load-sharing is enabled by default. However, it has no effect unless IP routing and OSPF are enabled.*

< 1 - 4 >

Specifies the maximum number of equal-cost next hop paths the router allows. (Range: 2 - 4; Default: 4)

Displaying the Current IP Load-Sharing Configuration

Use the **show running** command to view the currently active IP load-sharing configuration, and **show config** to view the IP load-sharing configuration in the startup-config file. (While in its default configuration, IP load-sharing does not appear in the command output.) If IP load sharing is configured with non-default settings (disabled or configured for either two or three equal-cost next-

hop paths), then the current settings are displayed in the command output.

```
ProCurve(config)# show running
Running configuration:
; J8697A Configuration Editor; Created on
release #K.11.00
hostname "ProCurve"
module 1 type J8702A
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24
    ip address dhcp-bootp
    exit
[ip_load-sharing_3]
access-controller vlan-base 2000
```

Indicates a non-default IP load-sharing configuration allowing three equal-cost next-hop paths for routed traffic with different subnet destinations. If the routing switch is configured with the default IP load-sharing configuration, IP load-sharing does not appear in the **show config** or **show running** command output.

Figure 5-52. Displaying a Non-Default IP Load-Sharing Configuration

Configuring IRDP

The ICMP Router Discovery Protocol (IRDP) is used by ProCurve routing switches to advertise the IP addresses of their router interfaces to directly attached hosts. IRDP is disabled by default. You can enable the feature on a global basis or on an individual VLAN interface basis.

When IRDP is enabled, the routing switch periodically sends Router Advertisement messages out the IP interfaces on which the feature is enabled. The messages advertise the routing switch's IP addresses to directly attached hosts who listen for the messages. In addition, hosts can be configured to query the routing switch for the information by sending Router Solicitation messages.

Some types of hosts use the Router Solicitation messages to discover their default gateway. When IRDP is enabled on the ProCurve routing switch, the routing switch responds to the Router Solicitation messages. Some clients interpret this response to mean that the routing switch is the default gateway. If another router is actually the default gateway for these clients, leave IRDP disabled on the ProCurve routing switch.

IRDP uses the following parameters. If you enable IRDP on individual VLAN interfaces, you can configure these parameters on an individual VLAN interface basis.

- **Packet type** - The routing switch can send Router Advertisement messages as IP broadcasts or as IP multicasts addressed to IP multicast group 224.0.0.1. The default packet type is IP broadcast.
- **Hold time** - Each Router Advertisement message contains a hold time value. This value specifies the maximum amount of time the host should consider an advertisement to be valid until a newer advertisement arrives. When a new advertisement arrives, the hold time is reset. The hold time is always longer than the maximum advertisement interval. Therefore, if the hold time for an advertisement expires, the host can reasonably conclude that the router interface that sent the advertisement is no longer available. The default hold time is three times the maximum message interval.
- **Maximum message interval and minimum message interval** - when IRDP is enabled, the routing switch sends the Router Advertisement messages every 450-600 seconds by default. The time within this interval that the routing switch selects is random for each message and is not affected by traffic loads or other network factors. The random interval minimizes the probability that a host will receive Router Advertisement

messages from other routers at the same time. The interval on each IRDP-enabled routing switch interface is independent of the interval on other IRDP-enabled interfaces. The default maximum message interval is 600 seconds. The default minimum message interval is 450 seconds.

- **Preference** - If a host receives multiple Router Advertisement messages from different routers, the host selects the router that send the message with the highest preference as the default gateway. The preference can be a number from -4294967296 to 4294967295. The default is 0.

Enabling IRDP Globally

To enable IRDP globally, enter the following command:

```
ProCurve(config)# ip irdp
```

This command enables IRDP on the IP interfaces on all ports. Each port uses the default values for the IRDP parameters.

Enabling IRDP on an Individual VLAN Interface

To enable IRDP on an individual VLAN interface and configure IRDP parameters, enter commands such as the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip irdp maxadvertinterval 400
```

This example shows how to enable IRDP on a specific interface (VLAN 1) and change the maximum advertisement interval for Router Advertisement messages to 400 seconds.

Syntax: [no] ip irdp [broadcast | multicast] [holdtime <seconds>] [maxadvertinterval <seconds>] [minadvertinterval <seconds>] [preference <number>]

- **broadcast | multicast** - This parameter specifies the packet type the routing switch uses to send the Router Advertisement.
 - **broadcast** - The routing switch sends Router Advertisements as IP broadcasts.
 - **multicast** - The routing switch sends Router Advertisements as multicast packets addressed to IP multicast group 224.0.0.1. This is the default.
- **holdtime <seconds >** - This parameter specifies how long a host that receives a Router Advertisement from the routing switch should consider the advertisement to be valid. When a host receives a new Router Advertisement message from the routing switch, the host resets the hold time

for the routing switch to the hold time specified in the new advertisement. If the hold time of an advertisement expires, the host discards the advertisement, concluding that the router interface that sent the advertisement is no longer available. The value must be greater than the value of the `maxadvertinterval` parameter and cannot be greater than 9000. The default is three times the value of the `maxadvertinterval` parameter.

- **maxadvertinterval** - This parameter specifies the maximum amount of time the routing switch waits between sending Router Advertisements. You can specify a value from 1 to the current value of the `holdtime` parameter. The default is 600 seconds.
- **minadvertinterval** - This parameter specifies the minimum amount of time the routing switch can wait between sending Router Advertisements. The default is three-fourths (0.75) the value of the `maxadvertinterval` parameter. If you change the `maxadvertinterval` parameter, the software automatically adjusts the `minadvertinterval` parameter to be three-fourths the new value of the `maxadvertinterval` parameter. If you want to override the automatically configured value, you can specify an interval from 1 to the current value of the `maxadvertinterval` parameter.
- **preference < number >** - This parameter specifies the IRDP preference level of this routing switch. If a host receives Router Advertisements from multiple routers, the host selects the router interface that sent the message with the highest preference as the host's default gateway. The valid range is -4294967296 to 4294967295. The default is 0.

Displaying IRDP Information

To display IRDP information, enter `show ip irdp` from any CLI level.

```
ProCurve# show ip irdp
Status and Counters - ICMP Router Discovery Protocol

Global Status : Disabled

VLAN Name      Status   Advertising   Min int   Max int   Holdtime   Preference
-----
Address        (sec)    (sec)        (sec)
-----
DEFAULT_VLAN   Enabled  multicast     450      600      1800      0
VLAN20         Enabled  multicast     450      600      1800      0
VLAN30         Enabled  multicast     450      600      1800      0
```

Figure 5-53. Example of Output for Show IP IRDP

Configuring DHCP Relay

Overview

The Dynamic Host Configuration Protocol (DHCP) is used for configuring hosts with IP address and other configuration parameters without user intervention. The protocol is composed of three components:

- DHCP client
- DHCP server
- DHCP relay agent

The DHCP client sends broadcast request packets to the network; the DHCP servers respond with broadcast packets that offer IP parameters, such as an IP address for the client. After the client chooses the IP parameters, communication between the client and server is by unicast packets.

ProCurve routing switches provide the DHCP relay agent to enable communication from a DHCP server to DHCP clients on subnets other than the one the server resides on. The DHCP relay agent transfers DHCP messages from DHCP clients located on a subnet without a DHCP server to other subnets. It also relays answers from DHCP servers to DHCP clients.

The DHCP relay agent is transparent to both the client and the server. Neither side is aware of the communications that pass through the DHCP relay agent. As DHCP clients broadcast requests, the DHCP relay agent receives the packets and forwards them to the DHCP server. During this process, the DHCP relay agent increases the hop count by one before forwarding the DHCP message to the server. A DHCP server includes the hop count from the DHCP request that it receives in the response that it returns to the client.

DHCP Packet Forwarding

The DHCP relay agent on the routing switch forwards DHCP client packets to all DHCP servers that are configured in the table administrated for each VLAN.

Unicast Forwarding

The packets are forwarded using unicast forwarding if the IP address of the DHCP server is a specific host address. The DHCP relay agent sets the destination IP address of the packet to the IP address of the DHCP server and forwards the message.

Broadcast Forwarding

The packets are forwarded using broadcast forwarding if the IP address of the DHCP server is a subnet address or IP broadcast address (255.255.255.255). The DHCP relay agent sets the DHCP server IP address to broadcast IP address and will be forwarded to all VLANs with configured IP interfaces (except the source VLAN).

Prerequisites for DHCP Relay Operation

For the DHCP Relay agent to work on the switch, you must complete the following steps:

1. Enable DHCP Relay on the routing switch (the default setting).
2. Ensure that a DHCP server is servicing the routing switch.
3. Enable IP Routing on the routing switch.
4. Ensure that there is a route from the DHCP server to the routing switch and back.
5. Configure one or more IP helper addresses for specified VLANs to forward DHCP requests to DHCP servers on other subnets.

Enabling DHCP Relay

The DHCP Relay function is enabled by default on a ProCurve routing switch. However, if DHCP has been disabled, you can re-enable it by entering the following command at the global configuration level:

```
ProCurve(config)# dhcp-relay
```

To disable the DHCP Relay function, enter the **no** form of the command:

```
ProCurve(config)# no dhcp-relay
```

Configuring an IP Helper Address

To add the IP address of a DHCP server for a specified VLAN on a routing switch, enter the **ip helper-address** command at the VLAN configuration level as in the following example:

```
ProCurve(config)# vlan 1  
ProCurve(vlan-1)# ip helper-address <ip-addr>
```

To remove the DHCP server helper address, enter the **no** form of the command:

```
ProCurve(vlan-1)# no ip helper-address <ip-addr>
```

Operating Notes

- You can configure up to 4000 IP helper addresses on a routing switch. The helper addresses are shared between the DHCP relay agent and UDP forwarder feature.
- A maximum of sixteen IP helper addresses is supported in each VLAN.

Hop Count in DHCP Requests

When a DHCP client broadcasts requests, the DHCP relay agent in the routing switch receives the packets and forwards them to the DHCP server (on a different subnet, if necessary). During this process the DHCP relay agent increments the hop count before forwarding DHCP packets to the server. The DHCP server, in turn, includes the hop count from the received DHCP request in the response sent back to a DHCP client.

As a result, the DHCP client receives a non-zero hop count in the DHCP response packet. Because some legacy DHCP/BootP clients discard DHCP responses which contain a hop count greater than one, they may fail to boot up properly. Although this behavior is in compliance with RFC 1542, it prevents a legacy DHCP/BootP client from being automatically configured with a network IP address.

Disabling the Hop Count in DHCP Requests

To disable the default behavior of a DHCP relay agent so that the hop count in a DHCP client request is not increased by one at each hop when it is forwarded to a DHCP server, enter the **no dhcp-relay hop-count-increment** command at the global configuration level:

```
ProCurve(config)# no dhcp-relay hop-count-increment
```

To reset the default function which increases the hop count in each DHCP request forwarded to a DHCP server, enter the following command:

```
ProCurve(config)# dhcp-relay hop-count-increment
```

Operating Notes

- By default, the DHCP relay agent increases the hop count in each DHCP request by one. You must enter the **no dhcp-relay hop-count-increment** command to disable this function.
- You enter the **no dhcp-relay hop-count-increment** command at the global configuration level. The command is applied to all interfaces on the routing switch that are configured to forward DHCP requests.
- This DHCP Relay enhancement only applies to DHCP requests forwarded to a DHCP server. The server does not change the hop count included in the DHCP response sent to DHCP clients.
- When you disable or re-enable the DHCP hop count function, no other behavior of the relay agent is affected.
- You can configure the DHCP Relay hop count function only from the CLI; you cannot configure this software feature from the drop-down menus.
- A new MIB variable, `hpDhcpRelayHopCount`, is introduced to support SNMP management of the hop count increment by the DHCP relay agent in a switch.

Verifying the DHCP Relay Configuration

Displaying the DHCP Relay Setting

Use the **show config** command (or **show running** for the running-config file) to display the current DHCP Relay setting.

Note

The DHCP relay and hop count increment settings appear in the **show config** command output only if the non-default values are configured. For more information about the DHCP hop count increment, see “Hop Count in DHCP Requests” on page 5-125.

```
ProCurve# show config
Startup configuration:
; J8697A Configuration Editor; Created on release #K.11.00
hostname "ProCurve"
cdp run
module 1 type J8702A
ip default-gateway 18.30.240.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1
  ip address 18.30.240.180 255.255.248.0
  no untagged A2-A24
  exit
no dhcp-relay
no dhcp-relay hop-count-increment
```

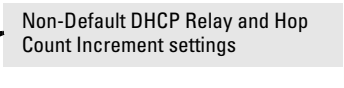


Figure 5-54. Displaying Startup Configuration with DHCP Relay and Hop Count Increment Disabled

Displaying DHCP Helper Addresses

To display the list of currently configured IP Helper addresses for a specified VLAN on the switch, enter the **show ip helper-address vlan** command.

Syntax: show ip helper-address [vlan <vlan-id>]

*Displays the IP helper addresses of DHCP servers configured for all static VLANs in the switch or on a specified VLAN, regardless of whether the DHCP Relay feature is enabled. The **vlan <vlan-id>** parameter specifies a VLAN ID number.*

The following command lists the currently configured IP Helper addresses for VLAN 1.

```
ProCurve(config)# show ip helper-address vlan 1
IP Helper Addresses
IP Helper Address
-----
10.28.227.97
10.29.227.53
```

Figure 5-55. Displaying IP Helper Addresses

Displaying the Hop Count Setting

To verify the current setting for increasing the hop count in DHCP requests, enter the **show dhcp-relay** command. Note that the current setting is displayed next to DHCP Request Hop Count Increment.

```
ProCurve# show dhcp-relay

Status and Counters - DHCP Relay Agent
DHCP Relay Agent Enabled           : Yes
DHCP Request Hop Count Increment: Disabled
Option 82 Handle Policy           : Replace
Remote ID                          : MAC Address

Client Requests      Server Responses
Valid      Dropped   Valid      Dropped
-----
1425          2       1425        0
```

DHCP Option 82

Option 82 is called the Relay Agent Information option and is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server. Servers recognizing the Relay Agent Information option may use the information to implement IP address or other parameter assignment policies. The DHCP Server echoes the option back verbatim to the relay agent in server-to-client replies, and the relay agent strips the option before forwarding the reply to the client.

The “Relay Agent Information” option is organized as a single DHCP option that contains one or more “sub-options” that convey information known by the relay agent. The initial sub-options are defined for a relay agent that is co-located in a public circuit access unit. These include a “circuit ID” for the incoming circuit, and a “remote ID” which provides a trusted identifier for the remote high-speed modem.

The routing switch can operate as a DHCP relay agent to enable communication between a client and a DHCP server on a different subnet. Without Option 82, DHCP operation modifies client IP address request packets to the extent needed to forward the packets to a DHCP server. Option 82 enhances this operation by enabling the routing switch to append an *Option 82 field* to such client requests. This field includes two suboptions for identifying the routing switch (by MAC address or IP address) and the routing switch port the client is using to access the network. A DHCP server with Option 82 capability can read the appended field and use this data as criteria for selecting the IP

addressing it will return to the client through the usual DHCP server response packet. This operation provides several advantages over DHCP without Option 82:

- An Option 82 DHCP server can use a relay agent's identity and client source port information to administer IP addressing policies based on client and relay agent location within the network, regardless of whether the relay agent is the client's primary relay agent or a secondary agent.
- A routing switch operating as a primary Option 82 relay agent for DHCP clients requesting an IP address can enhance network access protection by blocking attempts to use an invalid Option 82 field to imitate an authorized client, or by blocking attempts to use response packets with missing or invalid Option 82 suboptions to imitate valid response packets from an authorized DHCP server.
- An Option 82 relay agent can also eliminate unnecessary broadcast traffic by forwarding an Option 82 DHCP server response only to the port on which the requesting client is connected, instead of broadcasting the DHCP response to all ports on the VLAN.

Note

The routing switch's DHCP Relay Information (Option 82) feature can be used in networks where the DHCP server(s) are compliant with RFC 3046 Option 82 operation. DHCP Servers that are not compliant with Option 82 operation ignore Option 82 fields. For information on configuring an Option 82 DHCP server, refer to the documentation provided with the server application.

Some client applications can append an Option 82 field to their DHCP requests. Refer to the documentation provided for your client application.

It is not necessary for all relay agents on the path between a DHCP client and the server to support Option 82, and a relay agent without Option 82 should forward DHCP packets regardless of whether they include Option 82 fields. However, Option 82 relay agents should be positioned at the DHCP policy boundaries in a network to provide maximum support and security for the IP addressing policies configured in the server.

Option 82 Server Support

To apply DHCP Option 82, the routing switch must operate in conjunction with a server that supports Option 82. (DHCP servers that do not support Option 82 typically ignore Option 82 fields.) Also, the routing switch applies Option 82 functionality only to client request packets being *routed* to a DHCP server. DHCP relay with Option 82 does not apply to *switched* (non-routed) client requests.

For information on configuring policies on a server running DHCP Option 82, refer to the documentation provided for that application.

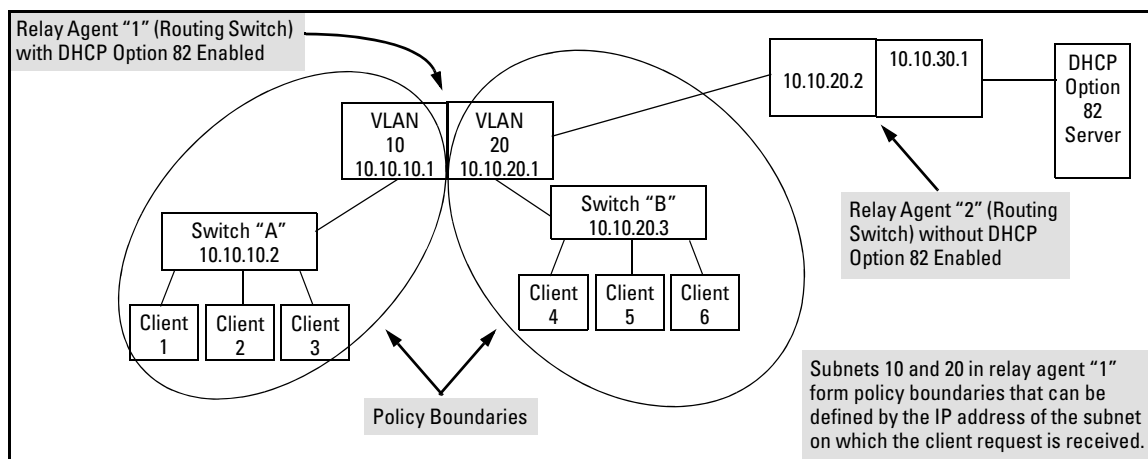


Figure 5-56. Example of a DHCP Option 82 Application

Terminology

Circuit ID: In Option 82 applications, the number of the port through which the routing switch receives a DHCP client request. On ProCurve fixed-port switches, the Circuit ID of a given port corresponds to the port number appearing on the front of the switch for that port. On ProCurve chassis switches, the port number for a given port corresponds to the internal if Index number for that port. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Circuit ID, refer to “Circuit ID” in the list on page 5-133.)

DHCP Policy Boundary: For Option 82 applications, an area of a network as defined by connection to a given routing switch or subnet and/or a specific port belonging to the routing switch or subnet.

DHCP relay agent: See Relay Agent.

Forwarding Policy: The Option 82 method the routing switch uses to process incoming client DHCP requests. For a given inbound DHCP client request, the forwarding policy determines whether the routing switch will add Option 82 information, replace existing Option 82 information, or leave any existing information unchanged. The policy also determines whether the routing switch will forward the client request toward a DHCP server or drop the

request. For a DHCP server response to an Option 82 client request, the routing switch can optionally perform a validation check to determine whether to forward or drop the response. Each Option 82 relay agent in the path between a DHCP client and an Option 82 DHCP server can be configured with a unique forwarding policy, which enhances DHCP policy control over discrete areas of a network.

Primary Relay Agent: In the path between a DHCP client and a DHCP server, the first routing switch (configured to support DHCP operation) that a client DHCP request encounters in the path from the client to a DHCP server.

Relay Agent: A routing switch that is configured to support DHCP operation.

Remote ID: In Option 82 applications on ProCurve switches, either the MAC address of a relay agent or the IP address of a VLAN or subnet configured on a relay agent or the (optional) Management VLAN configured on a relay agent. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Remote ID, refer to “Remote ID” in the bulleted list on page 5-132.)

Secondary Relay Agent: In the path between a DHCP client and a DHCP server, any routing switch (configured to support DHCP operation) other than the primary relay agent.

General DHCP Option 82 Requirements and Operation

Requirements. DHCP Option 82 operation is configured at the global config level and requires the following:

- IP routing enabled on the switch
- DHCP-Relay Option 82 enabled (global command level)
- routing switch access to an Option 82 DHCP server on a different subnet than the clients requesting DHCP Option 82 support
- one IP Helper address configured on each VLAN supporting DHCP clients

General DHCP-Relay Operation with Option 82. Typically, the first (primary) Option 82 relay agent to receive a client's DHCP request packet appends an Option 82 field to the packet and forwards it toward the DHCP server identified by the IP Helper address configured on the VLAN in which the client packet was received. Other, upstream relay agents used to forward the packet may append their own Option 82 fields, replace the Option 82 field(s) they find in the packet, forward the packet without adding another field, or drop the packet. (Intermediate next-hop routing switches without Option 82 capability can be used to forward—route—client request packets

with Option 82 fields.) Response packets from an Option 82 server are routed back to the primary relay agent (routing switch), and include an IP addressing assignment for the requesting client and an exact copy of the Option 82 data the server received with the client request. The relay agent strips off the Option 82 data and forwards the response packet out the port indicated in the response as the Circuit ID (client access port). Under certain validation conditions described later in this section, a relay agent detecting invalid Option 82 data in a response packet may drop the packet.

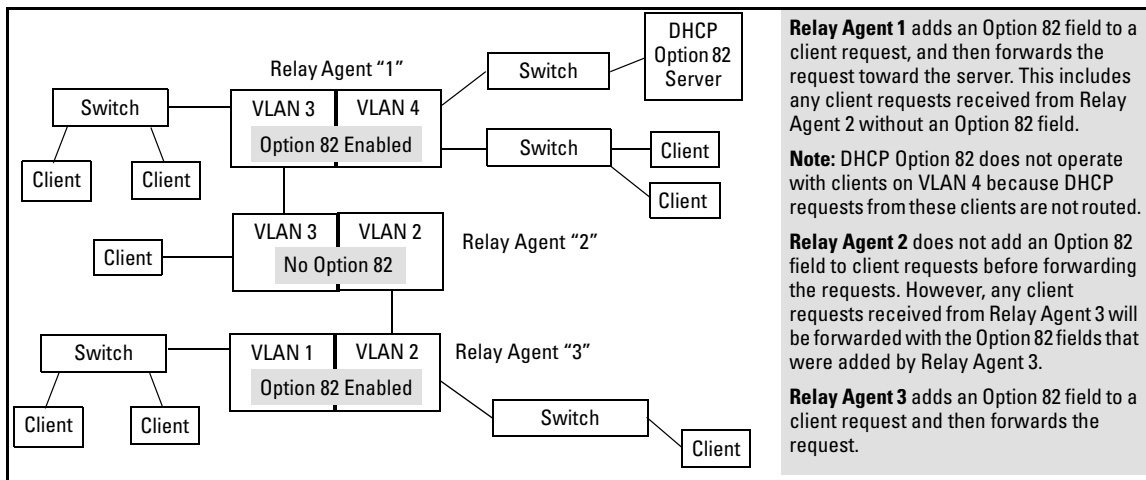


Figure 5-57. Example of DHCP Option 82 Operation in a Network with a Non-Compliant Relay Agent

Option 82 Field Content

The Remote ID and Circuit ID subfields comprise the Option 82 field a relay agent appends to client requests. A DHCP server configured to apply a different IP addressing policy to different areas of a network uses the values in these subfields to determine which DHCP policy to apply to a given client request.

- **Remote ID:** This configurable subfield identifies a policy area that comprises either the routing switch as a whole (by using the routing switch MAC address) or an individual VLAN configured on the routing switch (by using the IP address of the VLAN receiving the client request).
 - Use the IP address option if the server will apply different IP addressing policies to DHCP client requests from ports in different VLANs on the same routing switch.

- Use the Management VLAN option if a Management VLAN is configured and you want all DHCP clients on the routing switch to use the same IP address. (This is useful if you are applying the same IP addressing policy to DHCP client requests from ports in different VLANs on the same routing switch.) Configuring this option means the Management VLAN's IP address appears in the remote ID subfield of all DHCP requests originating with clients connected to the routing switch, regardless of the VLAN on which the requests originate.
- Use the MAC address option if, on a given routing switch, it does not matter to the DHCP server which VLAN is the source of a client request (that is, use the MAC address option if the IP addressing policies supported by the target DHCP server do not distinguish between client requests from ports in different VLANs in the same routing switch)

To view the MAC address for a given routing switch, execute the **show system-information** command in the CLI.

```
ProCurve(config)# show system-information

Status and Counters - General System Information

System Name       : ProCurve
System Contact    :
System Location   :

MAC Age Time (sec) : 300
Time Zone         : 0
Daylight Time Rule : None

Firmware revision : K.11.00   Base MAC Addr  : 00110a-a50c20
ROM Version       : K.11.00   Serial Number   : SG426NB048

Up Time          : 32 mins   Memory - Total  : 33,043,456
CPU Util (%)     : 4         Free           : 25,335,136

IP Mgmt - Pkts Rx : 0         Packet - Total  : 1998
          Pkts Tx : 0         Buffers - Free  : 1748
                                   Lowest  : 1741
                                   Missed   : 0
```

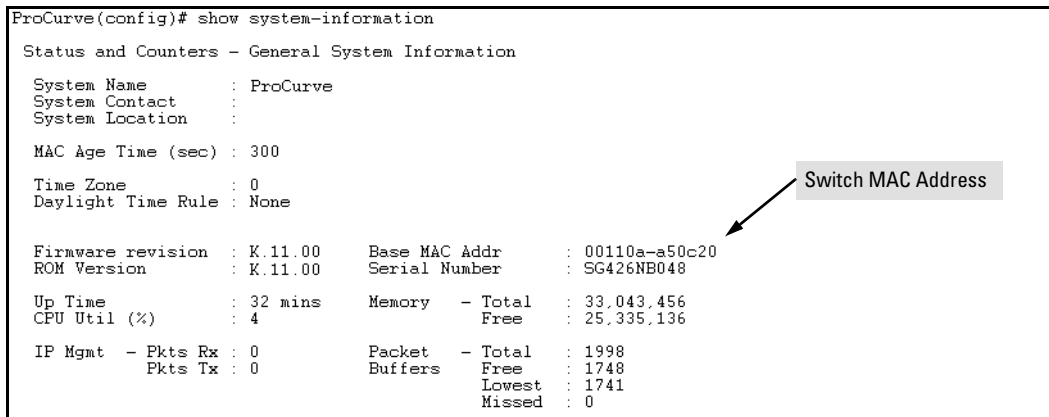


Figure 5-58. Using the CLI To View the Switch MAC Address

- **Circuit ID:** This nonconfigurable subfield identifies the port number of the physical port through which the routing switch received a given DHCP client request, and is necessary to identify if you want to configure an Option 82 DHCP server to use the Circuit ID to select a DHCP policy to assign to clients connected to the port. This number is the identity of the inbound port. On ProCurve fixed-port switches, the port number used for the Circuit ID is always the same as the physical port number shown on the front of the switch. On ProCurve chassis switches, where a dedicated, sequential block of internal port numbers are reserved for each slot, regardless of whether a slot is occupied, the circuit ID for a given port is

the sequential index number for that port position in the slot. (To view the Index number assignments for ports in the routing switch, use the **walkmib ifname** command.)

For example, the Circuit ID for port B11 on a ProCurve switch is “35”. (See Figure 5-59, below.)

```
ProCurve# walkmib ifname

ifName.1 = A1
ifName.2 = A2
ifName.3 = A3
ifName.4 = A4
ifName.25 = B1
ifName.26 = B2
ifName.27 = B3
ifName.28 = B4
ifName.29 = B5
ifName.30 = B6
ifName.31 = B7
ifName.32 = B8
ifName.33 = B9
ifName.34 = B10
ifName.35 = B11 ←
ifName.36 = B12
ifName.37 = B13
ifName.38 = B14
ifName.39 = B15
ifName.40 = B16
ifName.41 = B17
ifName.42 = B18
ifName.43 = B19

-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

In this example, the switch has a 4-port module installed in slot “A” and a 24-port module installed in slot “B”. Thus, the first port numbers in the listing are the Index numbers reserved for slot “A”. The first Index port number for slot “B” is “25”, and the Index port number for port B11 (and therefore the Circuit ID number) is “35”.

The Index (and Circuit ID) number for port B11 on the routing switch.

Figure 5-59. Using Walkmib To Determine the Circuit ID for a Port on a ProCurve Chassis

For example, suppose you wanted port 10 on a given relay agent to support no more than five DHCP clients simultaneously, you could configure the server to allow only five IP addressing assignments at any one time for the circuit ID (port) and remote ID (MAC address) corresponding to port 10 on the selected relay agent.

Similarly, if you wanted to define specific ranges of addresses for clients on different ports in the same VLAN, you could configure the server with the range of IP addresses allowed for each circuit ID (port) associated with the remote ID (IP address) for the selected VLAN.

Forwarding Policies

DHCP Option 82 on ProCurve switches offers four forwarding policies, with an optional validation of server responses for three of the policy types (**append**, **replace**, or **drop**).

Configuration Options for Managing DHCP Client Request Packets

Option 82 Configuration	DHCP Client Request Packet Inbound to the Routing Switch	
	Packet Has No Option 82 Field	Packet Includes an Option 82 Field
Append	Append an Option 82 Field	<p>Append allows the most detail in defining DHCP policy boundaries. For example, where the path from a client to the DHCP Option 82 server includes multiple relay agents with Option 82 capability, each relay agent can define a DHCP policy boundary and append its own Option 82 field to the client request packet. The server can then determine in detail the agent hops the packet took, and can be configured with a policy appropriate for any policy boundary on the path.</p> <p>Note: In networks with multiple relay agents between a client and an Option 82 server, append can be used only if the server supports multiple Option 82 fields in a client request. If the server supports only one Option 82 field in a request, consider using the keep option.</p>
Keep	Append an Option 82 Field	<p>If the relay agent receives a client request that already has one or more Option 82 fields, keep causes the relay agent to retain such fields and forward the request without adding another Option 82 field. But if the incoming client request does not already have any Option 82 fields, the relay agent appends an Option 82 field before forwarding the request. Some applications for keep include:</p> <ul style="list-style-type: none"> • The DHCP server does not support multiple Option 82 packets in a client request and there are multiple Option 82 relay agents in the path to the server. • The unusual case where DHCP clients in the network add their own Option 82 fields to their request packets and you do not want any additional fields added by relay agents. <p>This policy does not include the validate option (described in the next section) and allows forwarding of all server response packets arriving inbound on the routing switch (except those without a primary relay agent identifier.)</p>
Replace	Append an Option 82 Field	<p>Replace replaces any existing Option 82 fields from downstream relay agents (and/or the originating client) with an Option 82 field for the current relay agent. Some applications for replace include:</p> <ul style="list-style-type: none"> • The relay agent is located at a point in the network that is a DHCP policy boundary and you want to replace any Option 82 fields appended by downstream devices with an Option 82 field from the relay agent at the boundary. (This eliminates downstream Option 82 fields you do not want the server to use when determining which IP addressing policy to apply to a client request.) • In applications where the routing switch is the primary relay agent for clients that may append their own Option 82 field, you can use replace to delete these fields if you do not want them included in client requests reaching the server.

Option 82 Configuration	DHCP Client Request Packet Inbound to the Routing Switch	
	Packet Has No Option 82 Field	Packet Includes an Option 82 Field
Drop	Append an Option 82 Field	Drop causes the routing switch to drop an inbound client request with an Option 82 field already appended. If no Option 82 fields are present, drop causes the routing switch to add an Option 82 field and forward the request. As a general guideline, configure drop on relay agents at the edge of a network, where an inbound client request with an appended Option 82 field may be unauthorized, a security risk, or for some other reason, should not be allowed.

Multiple Option 82 Relay Agents in a Client Request Path

Where the client is one router hop away from the DHCP server, only the Option 82 field from the first (and only) relay agent is used to determine the policy boundary for the server response. Where there are multiple Option 82 router hops between the client and the server, you can use different configuration options on different relay agents to achieve the results you want. This includes configuring the relay agents so that the client request arrives at the server with either one Option 82 field or multiple fields. (Using multiple Option 82 fields assumes that the server supports multiple fields and is configured to assign IP addressing policies based on the content of multiple fields.)




Figure 5-60. Example Configured To Allow Only the Primary Relay Agent To Contribute an Option 82 Field

The above combination allows for detection and dropping of client requests with spurious Option 82 fields. If none are found, then the drop policy on the first relay agent adds an Option 82 field, which is then kept unchanged over the next two relay agent hops ("B" and "C"). The server can then enforce an IP addressing policy based on the Option 82 field generated by the edge relay agent ("A"). In this example, the DHCP policy boundary is at relay agent 1.




Figure 5-61. Example Configured To Allow Multiple Relay Agents To Contribute an Option 82 Field

This is an enhancement of the previous example. In this case, each hop for an accepted client request adds a new Option 82 field to the request. A DHCP server capable of using multiple Option 82 fields can be configured to use this approach to keep a more detailed control over leased IP addresses. In this example, the primary DHCP policy boundary is at relay agent “A”, but more global policy boundaries can exist at relay agents “B” and “C”.




Figure 5-62. Example Allowing Only an Upstream Relay Agent To Contribute an Option 82 Field

Like the first example, above, this configuration drops client requests with spurious Option 82 fields from clients on the edge relay agent. However, in this case, only the Option 82 field from the last relay agent is retained for use by the DHCP server. In this case the DHCP policy boundary is at relay agent “C”. In the previous two examples the boundary was with relay “A”.

Validation of Server Response Packets

A valid Option 82 server response to a client request packet includes a copy of the Option 82 field(s) the server received with the request. With validation disabled, most variations of Option 82 information are allowed, and the corresponding server response packets are forwarded.

Server response validation is an option you can specify when configuring Option 82 DHCP for **append**, **replace**, or **drop** operation. (Refer to “Forwarding Policies” on page 5-135.) Enabling validation on the routing switch can enhance protection against DHCP server responses that are either from untrusted sources or are carrying invalid Option 82 information.

With validation enabled, the relay agent applies stricter rules to variations in the Option 82 field(s) of incoming server responses to determine whether to forward the response to a downstream device or to drop the response due to invalid (or missing) Option 82 information. Table 5-20, below, describes relay agent management of DHCP server responses with optional validation enabled and disabled

Table 5-20. Relay Agent Management of DHCP Server Response Packets.

Response Packet Content	Option 82 Configuration	Validation Enabled on the Relay Agent	Validation Disabled (The Default)
Valid DHCP server response packet without an Option 82 field.	append, replace, or drop¹	Drop the server response packet.	Forward server response packet to a downstream device.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a <i>Remote ID</i> and <i>Circuit ID</i> combination that did not originate with the given relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a <i>Remote ID</i> that did not originate with the relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
All other server response packets ³	append, keep², replace, or drop¹	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.

¹Drop is the recommended choice because it protects against an unauthorized client inserting its own Option 82 field for an incoming request.

²A routing switch with DHCP Option 82 enabled with the **keep** option forwards all DHCP server response packets except those that are not valid for either Option 82 DHCP operation (compliant with RFC 3046) or DHCP operation without Option 82 support (compliant with RFC 2131).

³A routing switch with DHCP Option 82 enabled drops an inbound server response packet if the packet does not have any device identified as the primary relay agent (*giaddr* = null; refer to RFC 2131).

Multinetted VLANs

On a multinetted VLAN, each interface can form an Option 82 policy boundary within that VLAN if the routing switch is configured to use IP for the remote ID suboption. That is, if the routing switch is configured with IP as the remote ID option and a DHCP client request packet is received on a multinetted VLAN, the IP address used in the Option 82 field will identify the subnet on which the packet was received instead of the IP address for the VLAN. This enables an Option 82 DHCP server to support more narrowly defined DHCP policy boundaries instead of defining the boundaries at the VLAN or whole routing switch levels. If the MAC address option (the default) is configured instead, then the routing switch MAC address will be used regardless of which subnet was the source of the client request. (The MAC address is the same for all VLANs configured on the routing switch.)

Note that all request packets from DHCP clients in the different subnets in the VLAN must be able to reach any DHCP server identified by the IP Helper Address(es) configured on that VLAN.

Configuring Option 82

To configure DHCP Option 82 on a routing switch, enter the **dhcp-relay option 82** command.

Syntax: `dhcp-relay option 82 < append [validate] | replace [validate] | drop [validate] | keep > [ip | mac | mgmt-vlan]`

append: *Configures the switch to append an Option 82 field to the client DHCP packet. If the client packet has existing Option 82 field(s) assigned by another device, the new field is appended to the existing field(s).*

The appended Option 82 field includes the switch Circuit ID (inbound port number) associated with the client DHCP packet, and the switch Remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** option (below).*

replace: *Configures the switch to replace existing Option 82 field(s) in an inbound client DHCP packet with an Option 82 field for the switch.*

The replacement Option 82 field includes the switch circuit ID (inbound port number) associated with the client DHCP packet, and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** option (below).*

Syntax: dhcp-relay option 82 < append [validate] | replace [validate] | drop [validate] | keep >
[ip | mac | mgmt-vlan]

— Continued —

drop: Configures the routing switch to unconditionally drop any client DHCP packet received with existing Option 82 field(s). This means that such packets will not be forwarded. Use this option where access to the routing switch by untrusted clients is possible.

If the routing switch receives a client DHCP packet without an Option 82 field, it adds an Option 82 field to the client and forwards the packet. The added Option 82 field includes the switch circuit ID (inbound port number*) associated with the client DHCP packet, and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** option (below).

keep: For any client DHCP packet received with existing Option 82 field(s), configures the routing switch to forward the packet as-is, without replacing or adding to the existing Option 82 field(s).

[validate]: This option operates when the routing switch is configured with append, replace, or drop as a forwarding policy. With validate enabled, the routing switch applies stricter rules to an incoming Option 82 server response to determine whether to forward or drop the response. For more information, refer to "Validation of Server Response Packets" on page 5-137.

[ip | mac | mgmt-vlan]

This option specifies the remote ID suboption that the switch uses in Option 82 fields added or appended to DHCP client packets. The type of remote ID defines DHCP policy areas in the client requests sent to the DHCP server. If a remote ID suboption is not configured, then the routing switch defaults to the **mac** option. (Refer to "Option 82 Field Content" on page 5-132.)

ip: Specifies the IP address of the VLAN on which the client DHCP packet enters the switch.

mac: Specifies the routing switch's MAC address. (The MAC address used is the same MAC address that is assigned to all VLANs configured on the routing switch.) This is the default setting.

mgmt-vlan: Specifies the IP address of the (optional) Management VLAN configured on the routing switch. Requires that a Management VLAN is already configured on the switch. If the Management VLAN is multi-netted, then the primary IP address configured for the Management VLAN is used for the remote ID.

If you enter the **dhcp-relay option 82** command without specifying either **ip** or **mac**, the MAC address of the switch on which the packet was received from the client is configured as the remote ID. For information about the Remote ID values used in the Option 82 field appended to client requests, see "Option 82 Field Content" on page 5-132.

Example of Option 82 Configuration

In the routing switch shown below, option 82 has been configured with **mgmt-vlan** for the Remote ID.

```
ProCurve(config)# dhcp-relay option 82 append mgmt-vlan
```

The resulting effect on DHCP operation for clients X, Y, and Z is shown in table 5-21.

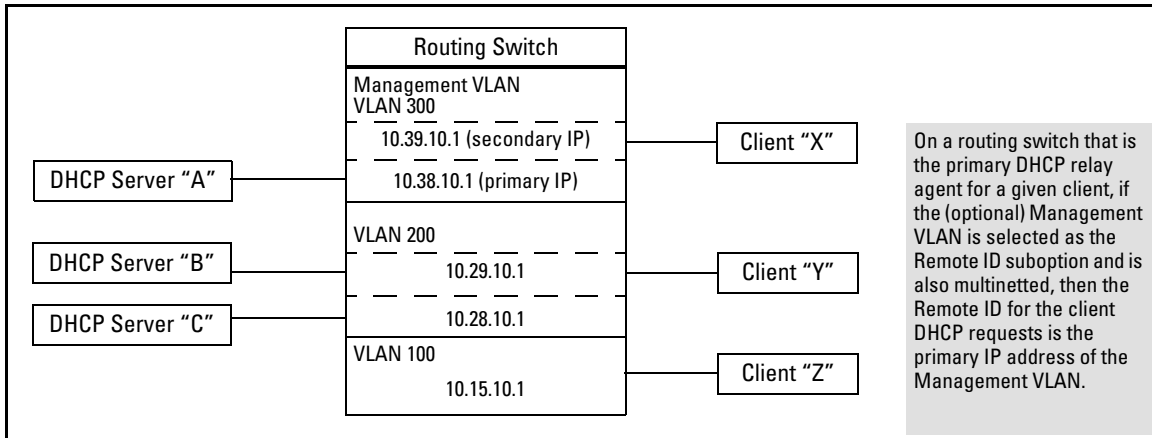


Figure 5-63. DHCP Option 82 When Using the Management VLAN as the Remote ID Suboption

Table 5-21. DHCP Operation for the Topology in Figure 5-63

Client	Remote ID	giaddr*	DHCP Server	
X	10.38.10.1	10.39.10.1	A only	If a DHCP client is in the Management VLAN, then its DHCP requests can go only to a DHCP server that is also in the Management VLAN. Routing to other VLANs is not allowed.
Y	10.38.10.1	10.29.10.1	B or C	Clients outside of the Management VLAN can send DHCP requests only to DHCP servers outside of the Management VLAN. Routing to the Management VLAN is not allowed.
Z	10.38.10.1	10.15.10.1	B or C	

*The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the giaddr (*gateway interface address*). This is the IP address of the VLAN on which the request packet was received from the client. For more information, refer to RFC 2131 and RFC 3046.

Operating Notes

- This implementation of DHCP relay with Option 82 complies with the following RFCs:
 - RFC 2131
 - RFC 3046
- Moving a client to a different port allows the client to continue operating as long as the port is a member of the same VLAN as the port through which the client received its IP address. However, rebooting the client after it moves to a different port can alter the IP addressing policy the client receives if the DHCP server is configured to provide different policies to clients accessing the network through different ports.
- The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the *giaddr* (gateway interface address). (That is, the *giaddr* is the IP address of the VLAN on which the request packet was received from the client.) For more information, refer to RFC 2131 and RFC 3046.
- DHCP request packets from multiple DHCP clients on the same relay agent port will be routed to the same DHCP server(s). Note that when using 802.1X on a switch, a port's VLAN membership may be changed by a RADIUS server responding to a client authentication request. In this case the DHCP server(s) accessible from the port may change if the VLAN assigned by the RADIUS server has different DHCP helper addresses than the VLAN used by unauthenticated clients.
- Where multiple DHCP servers are assigned to a VLAN, a DHCP client request cannot be directed to a specific server. Thus, where a given VLAN is configured for multiple DHCP servers, all of these servers should be configured with the same IP addressing policy.
- Where routing switch “A” is configured to insert its MAC address as the Remote ID in the Option 82 fields appended to DHCP client requests, and upstream DHCP servers use that MAC address as a policy boundary for assigning an IP addressing policy, then replacing switch “A” makes it necessary to reconfigure the upstream DHCP server(s) to recognize the MAC address of the replacement switch. This does not apply in the case where an upstream relay agent “B” is configured with **option 82 replace**, which removes the Option 82 field originally inserted by switch “A”.

- Relay agents without Option 82 can exist in the path between Option 82 relay agents and an Option 82 server. The agents without Option 82 will forward client requests and server responses without any effect on Option 82 fields in the packets.
- If the routing switch cannot add an Option 82 field to a client's DHCP request due to the message size exceeding the MTU (Maximum Transmission Unit) size, then the request is forwarded to the DHCP server without Option 82 data and an error message is logged in the switch's Event Log.
- Because routing is not allowed between the Management VLAN and other VLANs, a DHCP server must be available in the Management VLAN if clients in the Management VLAN require a DHCP server.
- If the Management VLAN IP address configuration changes after **mgmt-vlan** has been configured as the remote ID suboption, the routing switch dynamically adjusts to the new IP addressing for all future DHCP requests.
- The Management VLAN and all other VLANs on the routing switch use the same MAC address.

UDP Broadcast Forwarding

Overview

Some applications rely on client requests sent as limited IP broadcasts addressed to a UDP application port. If a server for the application receives such a broadcast, the server can reply to the client. Since typical router behavior, by default, does not allow broadcast forwarding, a client's UDP broadcast requests cannot reach a target server on a different subnet unless the router is configured to forward client UDP broadcasts to that server.

A switch with routing enabled includes optional per-VLAN UDP broadcast forwarding that allows up to 256 server and/or subnet entries on the switch (16 entries per-VLAN). If an entry for a particular UDP port number is configured on a VLAN and an inbound UDP broadcast packet with that port number is received on the VLAN, then the switch routes the packet to the appropriate subnet. (Each entry can designate either a single device or a single subnet. The switch ignores any entry that designates multiple subnets.)

Note

The number of UDP broadcast forwarding entries supported is affected by the number of IP helper addresses configured to support DHCP Relay. Refer to “Operating Notes for UDP Broadcast Forwarding” on page 5-149.

A UDP forwarding entry includes the desired UDP port number, and can be either an IP unicast address or an IP subnet broadcast address for the subnet the server is in. Thus, an incoming UDP packet carrying the configured port number will be:

- Forwarded to a specific host if a unicast server address is configured for that port number.
- Broadcast on the appropriate destination subnet if a subnet address is configured for that port number.

Note that a UDP forwarding entry for a particular UDP port number is always configured in a specific VLAN and applies only to client UDP broadcast requests received inbound on that VLAN. If the VLAN includes multiple subnets, then the entry applies to client broadcasts with that port number from any subnet in the VLAN.

For example, VLAN 1 (15.75.10.1) is configured to forward inbound UDP packets as shown in table 5-22:

Table 5-22. Example of a UDP Packet-Forwarding Environment

Interface	IP Address	Subnet Mask	Forwarding Address	UDP Port	Notes
VLAN 1	15.75.10.1	255.255.255.0	15.75.11.43	1188	Unicast address for forwarding inbound UDP packets with UDP port 1188 to a specific device on VLAN 2.
			15.75.11.255	1812	Broadcast address for forwarding inbound UDP packets with UDP port 1812 to any device in the 15.75.11.0 network.
			15.75.12.255	1813	Broadcast address for forwarding inbound UDP packets with UDP port 1813 to any device in the 15.75.12.0 network.
VLAN 2	15.75.11.1	255.255.255.0	<i>None</i>	<i>N/A</i>	Destination VLAN for UDP 1188 broadcasts from clients on VLAN 1. The device identified in the unicast forwarding address configured in VLAN 1 must be on this VLAN. Also the destination VLAN for UDP 1812 from clients on VLAN 1.
VLAN 3	15.75.12.1	255.255.255.0	<i>None</i>	<i>N/A</i>	Destination VLAN for UDP 1813 broadcasts from clients on VLAN 1.

Note

If an IP server or subnet entry is invalid, a switch will not try to forward UDP packets to the configured device or subnet address.

Subnet Masking for UDP Forwarding Addresses

The subnet mask for a UDP forwarding address is the same as the mask applied to the subnet on which the inbound UDP broadcast packet is received. To forward inbound UDP broadcast packets as limited broadcasts to other subnets, use the broadcast address that covers the subnet you want to reach. For example, if VLAN 1 has an IP address of 15.75.10.1/24 (15.75.10.1 255.255.255.0), then you can configure the following unicast and limited broadcast addresses for UDP packet forwarding to subnet 15.75.11.0:

Forwarding Destination Type	IP Address
UDP Unicast to a Single Device in the 15.75.11.0 Subnet	15.75.11.X
UDP Broadcast to Subnet 15.75.11.0	15.75.11.255

Configuring and Enabling UDP Broadcast Forwarding

To configure and enable UDP broadcast forwarding on the switch:

1. Enable routing.
2. Globally enable UDP broadcast forwarding.
3. On a per-VLAN basis, configure a forwarding address and UDP port type for each type of incoming UDP broadcast you want routed to other VLANs.

Globally Enabling UDP Broadcast Forwarding

Syntax [no] ip udp-bcast-forward

*Enables or disables UDP broadcast forwarding on the routing switch. Routing must be enabled before executing this command. Using the **no** form of this command disables any **ip forward-protocol udp** commands configured in VLANs on the switch. (Default: Disabled)*

Configuring UDP Broadcast Forwarding on Individual VLANs

This command routes an inbound UDP broadcast packet received from a client on the VLAN to the unicast or broadcast address configured for the UDP port type.

Syntax [no] ip forward-protocol udp < ip-address > < port-number | port-name >

*Used in a VLAN context to configure or remove a server or broadcast address and its associated UDP port number. You can configure a maximum of 16 **forward-protocol udp** assignments in a given VLAN. The switch allows a total of 256 **forward-protocol udp** assignments across all VLANs. You can configure UDP broadcast forwarding addresses regardless of whether UDP broadcast forwarding is globally enabled on the switch. However, the feature does not operate unless globally enabled.*

— Continued on the next page. —

— Continued from the preceding page. —

< ip-address >: This can be either of the following:

- The unicast address of a destination server on another subnet. For example: 15.75.10.43.
- The broadcast address of the subnet on which a destination server operates. For example, the following address directs broadcasts to All hosts in the 15.75.11.0 subnet: 15.75.11.255.

Note: The subnet mask for a forwarded UDP packet is the same as the subnet mask for the VLAN (or subnet on a multinetted VLAN) on which the UDP broadcast packet was received from a client.

< udp-port-# >: Any UDP port number corresponding to a UDP application supported on a device at the specified unicast address or in the subnet at the specified broadcast address. For more information on UDP port numbers, refer to “TCP/UDP Port Number Ranges” on page 5-149.

< port-name >: Allows use of common names for certain well-known UDP port numbers. You can type in the specific name instead of having to recall the corresponding number:

dns: Domain Name Service (53)

nntp: Network Time Protocol (123)

netbios-ns: NetBIOS Name Service (137)

netbios-dgm: NetBIOS Datagram Service (138)

radius: Remote Authentication Dial-In User Service (1812)

radius-old: Remote Authentication Dial-In User Service (1645)

rip: Routing Information Protocol (520)

snmp: Simple Network Management Protocol (161)

snmp-trap: Simple Network Management Protocol (162)

tftp: Trivial File Transfer Protocol (69)

timep: Time Protocol (37)

For example, the following command configures the routing switch to forward UDP broadcasts from a client on VLAN 1 for a time protocol server:

```
ProCurve(config)# ip forward-protocol udp 15.75.11.155  
timep
```

Displaying the Current IP Forward-Protocol Configuration

Syntax show ip forward-protocol [vlan < vid >]

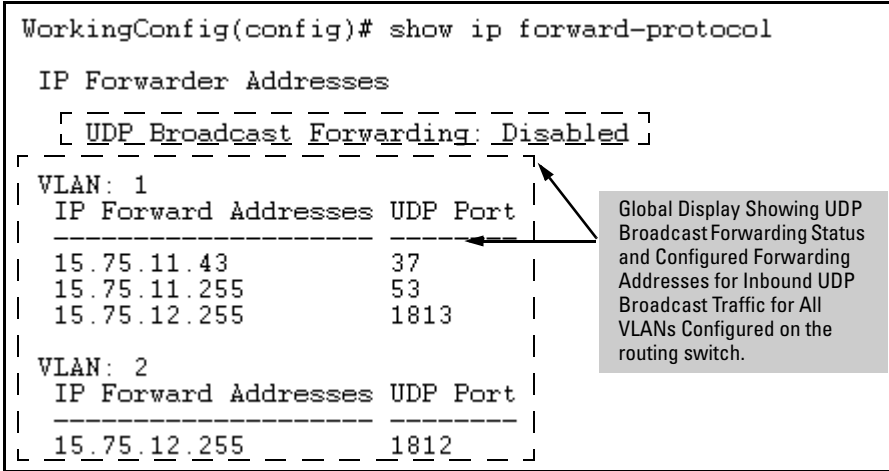
Displays the current status of UDP broadcast forwarding and lists the UDP forwarding address(es) configured on all static VLANs in the switch or on a specific VLAN.

```
WorkingConfig(config)# show ip forward-protocol

IP Forwarder Addresses
  [ UDP Broadcast Forwarding: Disabled ]

VLAN: 1
  IP Forward Addresses  UDP Port
  -----
  15.75.11.43           37
  15.75.11.255         53
  15.75.12.255         1813

VLAN: 2
  IP Forward Addresses  UDP Port
  -----
  15.75.12.255         1812
```



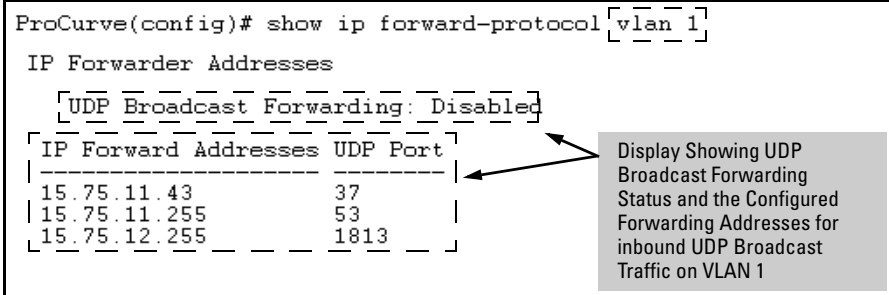
Global Display Showing UDP Broadcast Forwarding Status and Configured Forwarding Addresses for Inbound UDP Broadcast Traffic for All VLANs Configured on the routing switch.

Figure 5-64. Displaying Global IP Forward-Protocol Status and Configuration

```
ProCurve(config)# show ip forward-protocol [vlan 1]

IP Forwarder Addresses
  [ UDP Broadcast Forwarding: Disabled ]

  IP Forward Addresses  UDP Port
  -----
  15.75.11.43           37
  15.75.11.255         53
  15.75.12.255         1813
```



Display Showing UDP Broadcast Forwarding Status and the Configured Forwarding Addresses for inbound UDP Broadcast Traffic on VLAN 1

Figure 5-65. Displaying IP Forward-Protocol Status and Per-VLAN Configuration

Operating Notes for UDP Broadcast Forwarding

Maximum Number of Entries. The number of UDP broadcast entries and IP helper addresses combined can be up to 16 per VLAN, with an overall maximum of 2048 on the switch. (IP helper addresses are used with the switch’s DHCP Relay operation. For more information, refer to “Configuring DHCP Relay” on page 5-123.) For example, if VLAN 1 has 2 IP helper addresses configured, you can add up to 14 UDP forwarding entries in the same VLAN.

TCP/UDP Port Number Ranges. There are three ranges:

- Well-Known Ports: 0 - 1023
- Registered Ports: 1024 - 49151
- Dynamic and/or Private Ports: 49152 - 65535

For more information, including a listing of UDP/TCP port numbers, go to the *Internet Assigned Numbers Authority* (IANA) website at:

www.iana.org

Then click on:

Protocol Number Assignment Services

P (Under “Directory of General Assigned Numbers” heading)

Port Numbers

Messages Related to UDP Broadcast Forwarding

Message	Meaning
<code>udp-bcast-forward: IP Routing support must be enabled first.</code>	Appears in the CLI if an attempt to enable UDP broadcast forwarding has been made without IP routing being enabled first. Enable IP routing, then enable UDP broadcast forwarding.
<code>UDP broadcast forwarder feature enabled</code>	UDP broadcast forwarding has been globally enabled on the router. Appears in the Event Log and, if configured, in SNMP traps.
<code>UDP broadcast forwarder feature disabled</code>	UDP broadcast forwarding has been globally disabled on the routing switch. This action does not prevent you from configuring UDP broadcast forwarding addresses, but does prevent UDP broadcast forwarding operation. Appears in the Event Log and, if configured, in SNMP traps.
<code>UDP broadcast forwarder must be disabled first.</code>	Appears in the CLI if you attempt to disable routing while UDP forwarding is enabled on the switch.

