Release Notes *for* Version 05.0.84 of the HP ProCurve Routing Switch 9304M and 9308M Operating System

These release notes describe enhancements and fixes in HP ProCurve Routing Switch 9308M (J4138A) and 9304M (J4139A) software release 05.0.84, and are not included in the manual you received with the device. If you received this document with your HP ProCurve 9304M or 9308M routing switch, then release 05.0.84 is already installed in the device. Otherwise, you can download this release from HP's ProCurve website. To do so, go to **http://www.hp.com/go/procurve** and click on **Free Software Updates**. Then click on **Switches** and execute the download process for the file named **J4138084.exe**. For information on how to update your routing switch software, refer to chapter 3, "Updating Software Images and Configuration Files" in the *Installation and Configuration Guide* you received with the device.

Table of Contents

Summary of New Features in Release 05.0.84

Software release 05.0.84 contains the following enhancements:

- Border Gateway Protocol, version 4 (BGP4)

  IMPORTANT: This new feature required a change in the way the software handles default static routes. This change applies even if your do not use BGP4. See "Change in Treatment of Static Default Routes" on page 8.

- Local access control (support for up to 16 configurable user accounts with passwords)
- Loopback interfaces (router only)
- Enhanced supernetting (optimized CAM operation for three-level overlapping IP prefixes)
- Load sharing over default routes
- VLAN enhancement - a single router interface can span multiple port-based VLANs
- New command for increasing table sizes for OSPF Link-State Advertisements (LSAs)
- New command to display system defaults
- New commands to list and clear open Web management sessions
- New IP Multicast commands for configuring global IP Multicast parameters
- Enhancements to trunk port configuration to enable easier reconfiguration
- Page-display mode configurable from the User EXEC and Privileged EXEC CLI levels
- Source IP addresses in SNMP authentication failure traps
- New CLI command for displaying the running configuration file
- Configurable MAC table size
- TAB key displays command option list (same output as "?")
- Unneeded command (global-tag) removed from software

Software Image Files

The following table lists the boot and flash image files required to run software release 05.0.84.

Boot Image	Flash Image
BIB04500.bin BIB04600.bin	BIR05000.bin
Note: BIB04600.bin is required only for the J4141A 10/100 Management Mofule and the J4142 100Base FX module.	

Note

If you are adding the 24-port J4142A HP ProCurve 9300 100Base FX module to your HP 9304M or 9308M chassis and you want to be able use the boot monitor for TFTP access to the module, you must upgrade to boot code version BIB04600.bin. To view the boot image version, use the **show flash** command in the CLI.

Logging On

You can access an HP 9304M or 9308M through a direct serial connection to the CLI, through a Telnet connection to the CLI, or through a web browser.

- To access the CLI through a serial connection, attach a straight-through EIA/TIA DB-9 serial cable (M/F) to the management port on the chassis. Then use a terminal emulation application with the following settings to access the CLI:
  - Baud: 9600 bps
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None

  Note: A serial cable is included with your routing switch. If you prefer to build your own cable, see the pinout information in chapter 2, Installation, in your HP ProCurve Routing Switch 9304M and Routing Switch 9308M Installation and Configuration Guide.

- To access a routing switch's CLI through a Telnet connection, assign an IP address to a router interface attached to the network.

  a. Enter the enable command at the User EXEC level prompt (for example, `HP9304>`). Then press Enter. (If you are prompted for a password, enter your enable password.)

  b. Enter the configure terminal command at the Privileged EXEC level prompt (for example, `HP9300#`). Then press Enter.

  c. Enter the following command to access the Interface CONFIG level for the interface:

     int e <slot-num>/<port-num>, where **<slot num>** is the number of the slot in which the interface is installed, and **<port-num>** is the number of the port in the interface.

  d. Enter the following command to configure the IP address:

     ip address <ip address> <ip mask>

     Note: You also can enter ip address **<ip address /sub-net mask length>**.

  e. Enter the write memory command to save the configuration information to the device's flash memory.

- With the upgrade to release 5.0.84, you will not be able to access the Web management interface unless you configure the access options at the CLI level first. Therefore, if you did not previously have the access options configured, you will now need to do so in order to have web access.

Note

After upgrading to release 5.0.84, you must have access options configured at the CLI level before you can access the Web management interface.

The command for configuring the necessary access option is:

snmp-server community <private?> rw

where **<private?>** is the SNMP community name.

To access the Web management interface through a web browser, enter the device's IP address in the browser's Location or Address field, then press Enter. When the Login dialog is displayed, enter the default user name and password for read-only or read-write access:

- For read-write access, enter "set" in the User Name field and "private" in the Password field. (If this does not work, the read-write community name has been changed. Enter the read-write community name configured on the management server.)
- For read-only access, enter "get" in the User Name field and "public" in the Password field. (If this does not work, the read-only community name has been changed. Enter the read-only community name configured on the management server.)

Recovering from a Lost Password

By default, the CLI does not require passwords. However, if someone has configured a password for the device but the password has been lost, you can regain super-user access to the device using the following procedure.

Note

Recovery from a lost password requires direct access to the serial port and a system reset.

To recover from a lost password:

1. Start a CLI session over the serial interface to the device.
2. Reboot the device.
3. At the initial boot prompt at system startup, enter b to enter the boot monitor mode.
4. Enter no password at the prompt. (You cannot abbreviate this command.)
5. Enter boot system flash primary at the prompt. This command causes the device to bypass the system password check.
6. After the console prompt reappears, assign a new password.

Border Gateway Protocol (BGP) Version 4

Border Gateway Protocol (BGP) version 4 (often called BGP4), is the most widely used protocol for exchanging route information between autonomous systems in the Internet.

Chassis software release 05.0.84 provides an implementation of BGP4 based on the following RFCs:

- RFC 1771, "A Border Gateway Protocol 4"
- RFC 1745, "BGP4/IDRP for IP-OSPF Interaction"
- RFC 1997, "BGP Communities Attribute"

Note

BGP4 requires 32MB or greater of DRAM memory. Use the show version CLI command to view the amount of memory on the device. The last few lines of the resulting list will display the memory information.

This software release provides the following major features of the BGP4 protocol:

- Synchronization of BGP and Internal Gateway Protocol (IGP) routes

- Default route origination and learning
- Configurable administrative distance for default route
- Automatic summarization of IGP routes redistributed to BGP into Class-A, Class-B, or Class-C networks
- Route filtering and manipulation using route maps, address filters, AS-path filters, private-AS filtering option, and community filters
- Support for well-known and user-defined communities
- Loopback address support, including the ability to set the source address for TCP sessions (over which BGP neighbors exchange BGP messages) and source updates

See *Manual Supplement 2 for the HP ProCurve Routing Switch 9304M and 9308M* (available mid-August, 1999) for BGP4 configuration information. Go to HP's ProCurve website at **http://www.hp.com/go/procurve,** then click on **Technical Support**, then **Manual**.

Budgeting Memory for BGP4

BGP4 handles a very large number of routes and therefore requires a lot of memory. For example, in a typical configuration with just a single BGP4 neighbor, a BGP4 router many need to be able to hold up to 50,000 routes. Many configurations, especially those involving more than one neighbor, can require the router to hold even more routes.

Note

If you do not plan to use BGP4, you do not need to read this section.

HPdevices that support BGP4 can hold 30,000 routes by default. You can change the route table capacity on the HP 9304M and 9308M to a value from 10,000 - 70,000.

The software reserves memory for BGP4 route tables and other tables such as the IP route table and that memory cannot be used by other parts of the system. Therefore, the system will not operate properly unless all of the features you want to use have adequate memory. In addition, individual ports require memory, so you must ensure that the memory you allocate to route tables and other tables with configurable sizes leaves enough memory for the ports.

- For the HP 9304M and 9308M, you can configure memory to allow the maximum number of routes only if you also reconfigure memory to allow only one or two neighbors.
- To allow more than three neighbors, you must reduce the number of routes.

Change in Treatment of Default Static Routes

BGP4 uses an algorithm to evaluate multiple routes to the same destination and choose the best route. One of the parameters that the algorithm compares is the administrative distance of each route. The addition of BGP4 features in this software release required a change in the way the system handles default static routes.

IMPORTANT

This change affects the treatment of default static routes regardless of whether you use BGP4.

In previous software releases, if you configured a default static route, that route could not be overwritten, regardless of its administrative distance. Even if the routing switch learned a new default route with a lower distance, the new route did not replace the one you statically configured. In software release 05.0.84, it is possible for a default static route to be overwritten.

This can occur if the routing switch learns another default route that has a smaller administrative distance. For example, if the routing switch learns a default route through RIP or OSPF, and if the static route and the learned route both use their default administrative distances, the learned route has a lower distance and replaces the static route.

When comparing routes, the software gives preference to the smaller administrative distance. To prevent a static default route from being overwritten, reconfigure each static default route using the following command:

ip route <route-num> 0.0.0.0 0.0.0.0 <gateway> [<metric>] distance 100

This command reconfigures the static route and changes the administrative distance from the default 130 to 100. By changing the distance to 100, you give the route a lower distance than the default distances for RIP and OSPF routes. Thus, if the routing switch learns a default route through RIP and that route has the default RIP administrative distance 120, the static route (with distance 100) still has precedence and is not replaced. The **<metric>** is optional and defaults to 1. For example, to reconfigure static default route 1 with an administrative distance of 100, enter the following command:

HP9300(config)# ip route 1 0.0.0.0 0.0.0.0 209.157.22.1 distance 100

For reference, the following table lists the default administrative distances.

Note

These defaults apply even if you do not use BGP4.

Route Source	Default Administrative Distance
Directly-connected	0 (this value is not configurable)
External BGP (EBGP)	20
OSPF	110
RIP	120
Static	130 (applies to all static and default routes)
Internal BGP (IBGP)	200
Local BGP	200
Unknown	255 (makes the route unusable)

Local Access Control

Chassis software release 05.0.84 allows you to define up to 16 user accounts to control the following types of access to the HP 9304M and 9308M routing switches:

- Telnet access through the CLI
- Enable access through the EXEC level and CONFIG levels of the CLI
- Web-browser access through the Web management interface
- SNMP access through network management applications

When you configure a user account, you specify the CLI privilege level for the user (super-user, port-configuration, or read-only). For each type of access, you also configure an access-method list. An access-method lists specifies the authentication methods the device uses for access and the order in which the device tries the authentication methods.

The user accounts provide greater flexibility for controlling management access to HP 9304M and 9308M devices than the Enable passwords. You can continue to use the Enable passwords as a secondary means of access authentication or, if you prefer, you can choose not to use user accounts and instead continue to use only the Enable passwords. The local access feature is backward compatible with configuration files that contain Enable passwords. You also can use a RADIUS server for authentication.

For more information about user accounts, RADIUS authentication, access-method lists, and other access parameters and how to configure them, see the *Manual Supplement 2 for the HP ProCurve Routing Switch 9304M and 9308M*. Beginning with mid-August, 1999, this booklet will be included with routing switches shipped from the factory, and will also be available in PDF (portable document format) on HP's ProCurve website. To access this booklet, go to **http://www.hp.com/go/procurve,** then click on **Technical Support**, then **Manual**.

Loopback Interfaces

Prior to release 05.0.84 you could configure the following types of interfaces on the HP 9304M and 9308M:

- Ethernet interfaces, which are directly associated with physical ports
- Virtual interfaces (VEs), which are not directly associated with physical ports but are used for routing between port-based VLANs

You can configure either type of interface with a Layer 3 IP, IPX, or AppleTalk protocol address.

Software release 05.0.84 allows you to configure an additional type of interface, a loopback interface. A loopback interface is not associated with a specific port. Instead, a loopback interface is associated with an IP address.

Loopback interfaces are always up, regardless of the states of physical interfaces. For this reason, loopback addresses can be quite useful for management sessions with the routing switch. As long as at least one valid path to the routing switch exists, you can reach the device on one of its loopback interfaces.

You can add up to 24 IP addresses to each loopback interface.

Note

If you use a loopback interface to communicate with a BGP4 neighbor, the software always uses the lowest IP address configured on the loopback interface you specify. If you add a new IP address to a loopback interface that you have configured BGP4 to use when communicating with a BGP4 neighbor, and the new address is lower than the lowest address already configured on the loopback interface, you will need to reconfigure the neighbor entries in BGP4.

You can advertise loopback interfaces using RIP and OSPF. You also can redistribute the IP sub-nets of loopback interfaces into those protocols for advertisement. IP Multicast, VRRP, and SRP cannot operate on loopback interfaces. You cannot configure IPX or AppleTalk loopback interfaces.

Loopback interfaces also are quite useful for BGP4. In particular, when you configure BGP4 neighbor information on the routing switch, you can configure the routing switch to use a loopback interface's IP address to communicate with a specific BGP4 neighbor. By using a loopback interface's IP address, you can avoid problems that can be caused by frequent state changes on physical interfaces (sometimes called "interface flapping").

To configure a loopback interface, enter the following command at the global CLI level:

interface loopback <num>

where **<num>** is the interface number. You can enter a number from 1 - 8 for the 9304M and 9308M routing switches. Thus, you can configure up to four or eight loopback interfaces on the routing switch.

After you enter the above command, the CLI changes to the loopback interface level. At this level, you can configure an IP address on the interface.

The following example shows how to add loopback interface 1 with IP address 209.157.22.26 and a 16-bit network mask:

```
HP9300(config)# interface loopback 1
HP9300(config-lbif-1)# ip address 209.157.22.26/18
```

To configure the routing switch to use the loopback interface to communicate with a BGP4 neighbor, use a command such as the following at the BGP CONFIG level:

```
HP9300(config-bgp-router)# neighbor remote-as 100 209.157.22.12 update-source loopback 1
```

This command adds an IBGP (Interior Border Gateway Protocol) neighbor to Autonomous System (AS) 100, with IP address 209.157.22.12. The update-source loopback parameter configures the routing switch to use the IP address of a loopback interface (loopback interface 1 in this example) to communicate with the neighbor.

Note

Even if the BGP neighbor is local (an IBGP neighbor and thus located in the same Autonomous System as the HP routing switch), you need to use the `remote-as <IP addr>` parameter.

See *Manual Supplement 2 for the HP ProCurve Routing Switch 9304M and 9308M* (available mid-August, 1999) for BGP4 configuration information. Go to HP's ProCurve website at **http://www.hp.com/go/procurve,** then click on **Technical Support**, then **Manual**.

Enhanced Supernetting

Software release 05.0.84 contains an enhancement to the software that manages the routing switch's Content Addressable Memory (CAM). This enhancement provides greater CAM capacity, especially useful for environments with many IP addresses.

The look-up mechanism in software release 05.0.84 involves longest prefix match with up to three levels of overlapping prefixes. For the HP 9304M and 9308M, each individual module can have one or two distinct route caches from other modules. In other words, each module can have a separate route cache from every other module. The 10/100 modules each have one route cache (can be unique from other modules), and the Gigabit Ethernet modules have up to two distinct route caches each (one route cache per four ports).

This enhancement does not require any configuration changes.

Figure 1 (page 9) shows an example of the difference in CAM management between software release 05.0.84 and prior software. As shown in this simplified example, prior software registered the first route for a network (such as 10.10.20.1/24) as a network route.

Suppose an HP 9304M or 9308M routing switch has 3 routes in its IP routing table:

- 10.0.0.0/8 next-hop 5.5.5.1
- 10.10.0.0/16 next-hop 6.6.6.1
- 10.10.10.0/24 next-hop 7.7.7.1

Also suppose a host is sending packets to the following destinations:

- 10.10.10.1
- 10.10.10.2
- 10.10.10.3
- 10.10.20.1

- 10.10.20.2
- 10.10.20.3
- 10.20.0.1
- 10.20.0.2

Prior to release 05.0.84, the software added 10.10.10.0/24 to the CAM as a network route, then added all subsequent destinations that were not in 10.10.10.0/24 as host routes. This resulted in the following CAM entries:

- 10.10.10.0/24 -> 7.7.7.1
- 10.10.20.1/32 -> 6.6.6.1
- 10.10.20.2/32 -> 6.6.6.1
- 10.10.20.3/32 -> 6.6.6.1
- 10.20.0.1/32 -> 5.5.5.1
- 10.20.0.2/32 -> 5.5.5.1

Note that a total of six CAM entries were used.

In software release 05.0.84, the software adds the following entries to the CAM:

- 10.10.10.0/24 -> 7.7.7.1
- 10.10.0.0/16 -> 6.6.6.1
- 10.0.0.0/8 -> 5.5.5.1

Thus, only three CAM entries are used. Routes to hosts within the network routes in the CAM do not require additional CAM entries. Figure 1 shows the example described above.

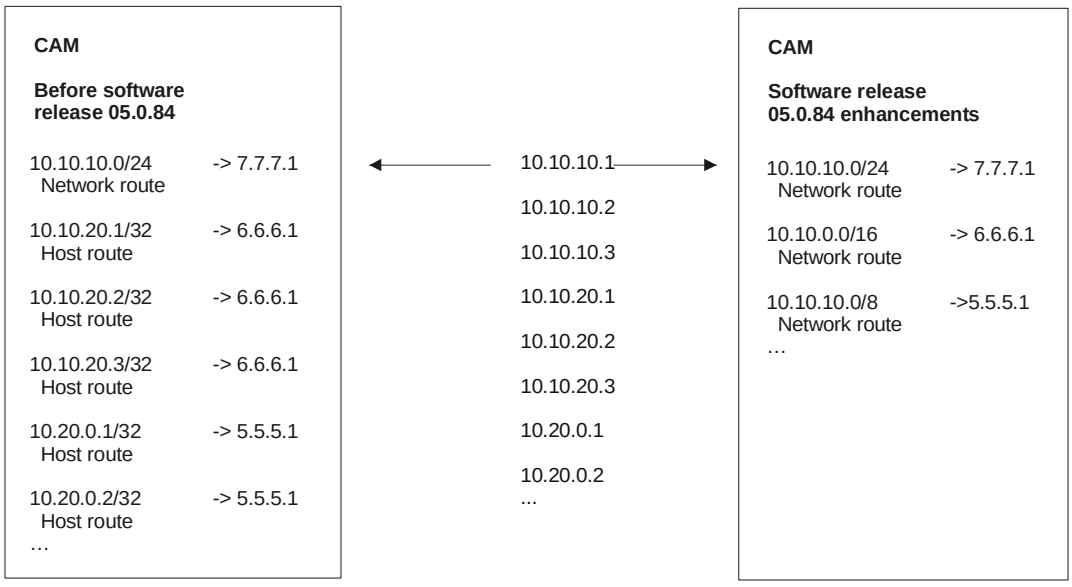

Figure 1. Example of CAM management changes

On HP routing switches running software release 05.0.84, the CAM can contain up to three levels of "supernet" host routes. In this example, the CAM contains 10.0.0.0/8, 10.10.0.0/16, and 10.10.10.0/24 network routes to accommodate all the possible host routes within those networks.

IP Load Sharing over Default Routes

In software releases earlier than release 05.0.84, the HP 9304M and 9308M could not load share across multiple paths on a default route. For example, if a default route configured in OSPF had multiple paths to next-hop Autonomous System Border Routers (ASBRs), the software could not load share among those paths. Figure 2 shows an example of this type of configuration.

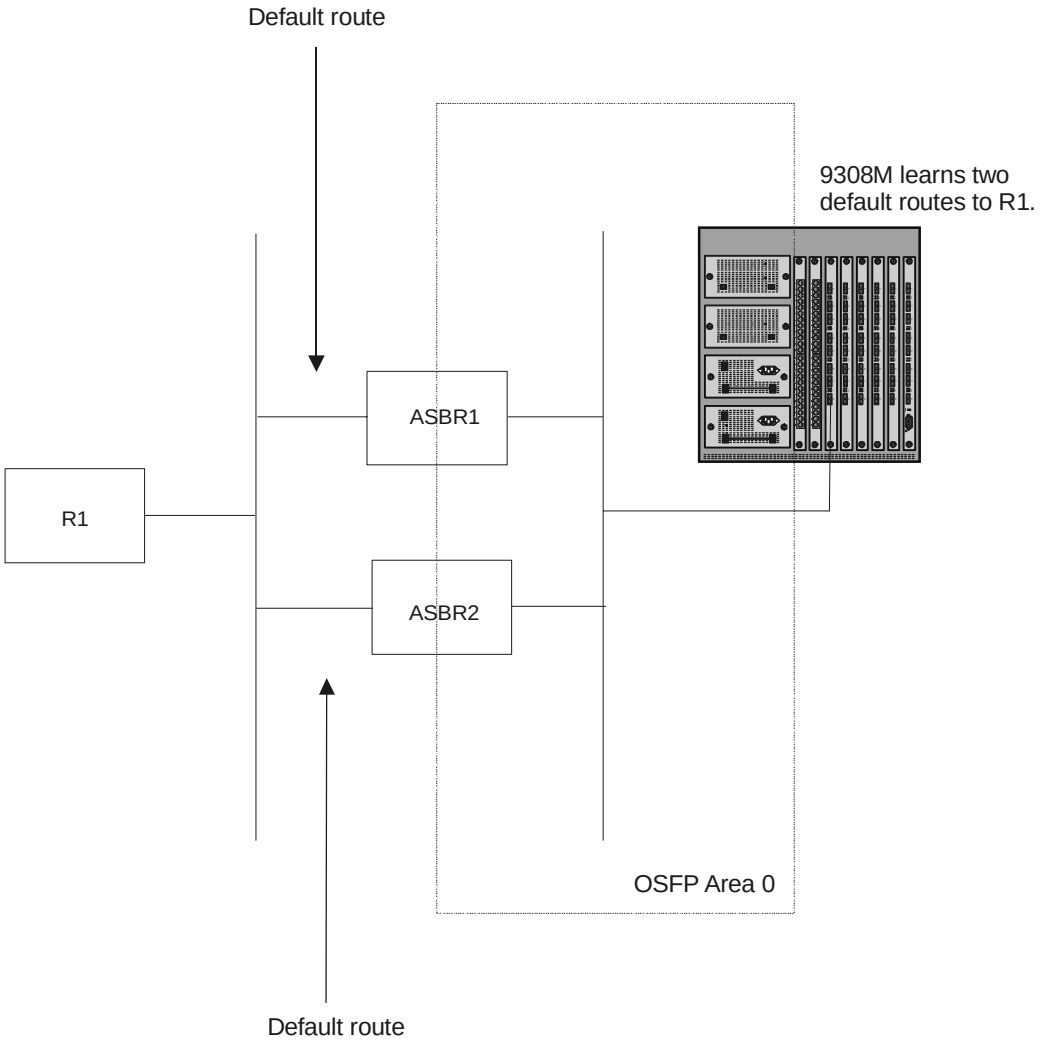

Figure 2. Default route with multiple paths to OSPF ASBRs

In this example, ASBR1 and ASBR2 each have a default route that they advertise through OSPF to the HP routing switch (an HP 9308M in this example). Prior to software release 05.0.84, the 9304M and 9308M could not load share between the default routes. Release 05.0.84 adds the capability to load share across multiple paths on default routes.

You can configure the HP 9304M and 9308M to load share among up to eight paths to the same destination.

Load sharing is disabled by default. The default number of paths for load sharing is four. To enable IP load sharing, enter the following command at the global CLI level:

ip load-sharing [<num>]

Where **<num>** specifies the maximum number of paths across which the device can load share. The range is 2 - 8 paths and the default is 4.

For more information about IP load sharing, see Chapter 8, "Configuring IP and IP RIP", in the *HP ProCurve Routing Switch 9304M and Routing Switch 9308M Installation and Configuration Guide* shipped with your routing switch. You can also access the manual on HP's ProCurve website by going to **http://www.hp.com/go/procurve**, clicking on **Technical Support**, and then clicking on **Manual**.

Multiple Port-Based VLANs with a Single IP Interface

For an HP 9304M or 9308M to route between port-based VLANs, you must add a virtual interface (VE) to each VLAN. Generally, you also configure a unique IP sub-net address on each virtual interface. For example, if you have three port-based VLANs, you add a virtual interface to each VLAN, then add a separate IP sub-net address to each virtual interface. The IP address on each of the virtual interfaces must be in a separate sub-net. The 9304M and 9308M route Layer 3 traffic between the sub-nets using the sub-net addresses. Figure 3 shows an example of this type of configuration.

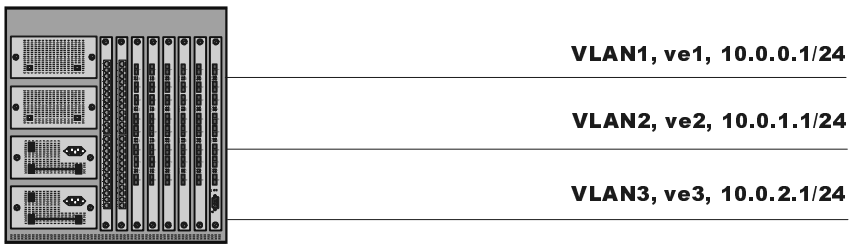

Figure 3. Multiple port-based VLANs with separate protocol addresses

As shown in this example, each VLAN has a separate IP sub-net address. If you need to conserve IP sub-net addresses, software release 05.0.84 adds the capability to configure multiple VLANs with the same IP sub-net address. For example, you can configure the VLANs in Figure 3 to use the same address, as shown in Figure 4.

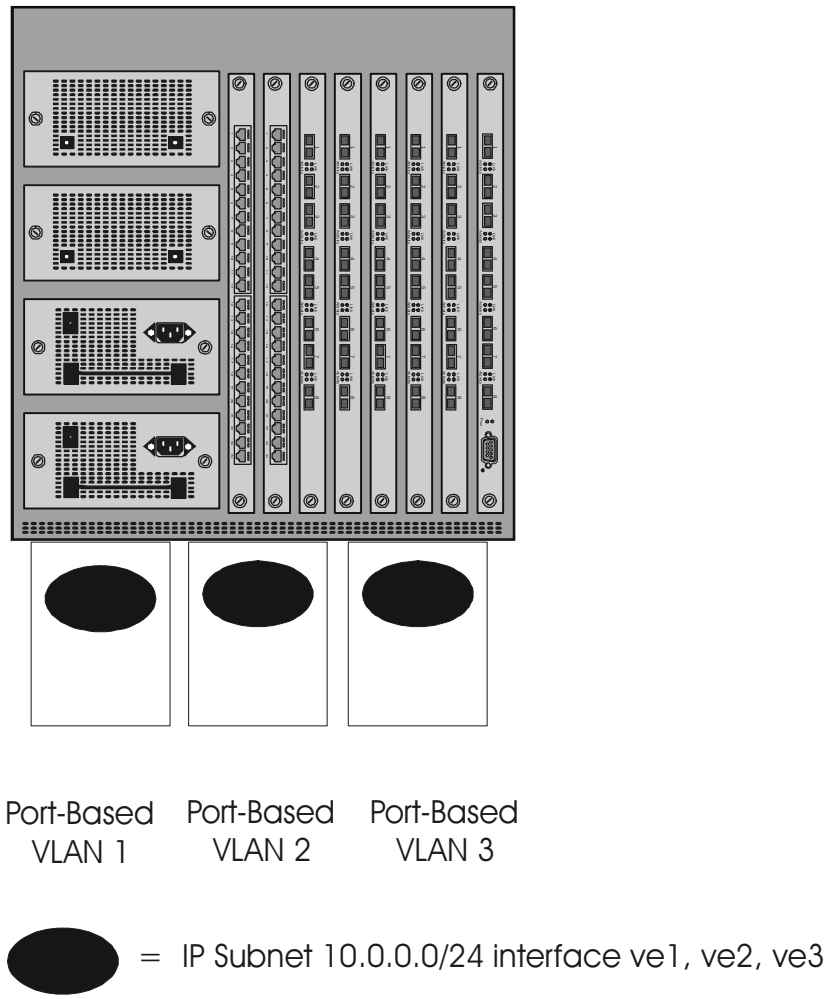

Figure 4. Multiple port-based VLANs with the same protocol address

Each VLAN still requires a separate virtual interface. However, all three VLANs now use the same IP sub-net address.

In addition to conserving IP sub-net addresses, this feature allows containment of Layer 2 broadcasts to segments within an IP sub-net. For ISP environments where the same IP sub-net is allocated to different customers, placing each customer in a separate VLAN allows all customers to share the IP sub-net address while at the same time being isolated from one another's Layer 2 broadcasts.

Note

You can provide redundancy to an IP sub-net address that contains multiple VLANs using a pair of HP routing switches configured for VRRP (Virtual Router Redundancy Protocol) or SRP (Standby Router Protocol). For more on VRRP, see *Manual Supplement 1 for the HP ProCurve Routing Switch 9304M and 9308M*, provided with HP routing switches shipped after mid-August, 1999, and also available on HP's ProCurve website. To access the supplement, go to **http://www.hp.com/go/procurve**. Click on **Technical Support**, then click on **Manual**.

The HP 9304M and 9308M perform proxy Address Resolution Protocol (ARP) for hosts that want to send IP traffic to hosts in other VLANs that are sharing the same IP sub-net address. If the source and destination hosts are in the same VLAN, the routing switch does not need to ARP.

- If a host attached to one VLAN sends an ARP message for the MAC address of a host in one of the other VLANs using the same IP sub-net address, the routing switch performs a proxy ARP on behalf of the other host. The routing switch then replies to the ARP by sending the virtual interface MAC address. The routing switch uses the same MAC address for all virtual interfaces.

  When the host that sent the ARP then sends a unicast packet addressed to the virtual interface's MAC address, the 9304M or 9308M switches the packet on Layer 3 to the destination host on the VLAN.

Note

If the routing switch's ARP table does not contain the requested host, it forwards the ARP request on Layer 2 to the same VLAN as the one that received the ARP request. Then the device sends an ARP for the destination to the other VLANs that are using the same IP sub-net address.

- If the destination is in the same VLAN as the source, the routing switch does not need to perform a proxy ARP.

To configure multiple VLANs to use the same IP sub-net address:

- Configure each VLAN, including adding tagged or untagged ports.
- Configure a separate virtual interface for each VLAN, but do not add an IP sub-net address to more than one of the virtual interfaces.
- Configure the virtual interfaces that do not have the IP sub-net address to "follow" the virtual interface that does have the address.

The CLI contains the following new command for configuring a virtual interface to follow another virtual interface:

```
ip follow ve <num>
```

Enter this command at the virtual interface configuration level. For example, to configure the VLANs shown in Figure 3 (page 11), you could enter the following commands.

```
HP9300(config)# vlan 1 by port
HP9300(config-vlan-1)# untag ethernet 1/1
HP9300(config-vlan-1)# tag ethernet 1/8
HP9300(config-vlan-1)# router-interface ve 1
```

The commands above configure port-based VLAN 1. The VLAN has one untagged port (1/1) and a tagged port (1/8). In this example, all three VLANs contain port 1/8 so the port must be tagged to allow the port to be in multiple VLANs. You can configure VLANs to share a Layer 3 protocol interface regardless of tagging. A combination of tagged and untagged ports is shown in this example to demonstrate that sharing the interface does not change other VLAN features.

Notice that each VLAN still requires a unique virtual interface (VE).

The following commands configure port-based VLANs 2 and 3.

```
HP9300(config-vlan-1)# vlan 2 by port
HP9300(config-vlan-2)# untag ethernet 1/2
HP9300(config-vlan-2)# tag ethernet 1/8
HP9300(config-vlan-2)# router-interface ve 2
HP9300(config-vlan-2)# vlan 3 by port
HP9300(config-vlan-3)# untag ethernet 1/5 to 1/6
HP9300(config-vlan-3)# tag ethernet 1/8
```

```
HP9300(config-vlan-3)# router-interface ve 3
```

The following commands configure an IP sub-net address on virtual interface 1.

```
HP9300(config-vlan-3)# interface ve 1
HP9300(config-vif-1)# ip address 10.0.0.1/24
```

The following commands configure virtual interfaces 2 and 3 to "follow" the IP sub-net address configured on virtual interface 1.

```
HP9300(config-vif-1)# interface ve 2
HP9300(config-vif-2)# ip follow ve 1
HP9300(config-vif-2)# interface ve 3
HP9300(config-vif-3)# ip follow ve 1
```

Note

Since virtual interfaces 2 and 3 do not have their own IP sub-net addresses but instead are "following" virtual interface 1's IP address, you still can configure an IPX or AppleTalk interface on virtual interfaces 2 and 3.

Increasing Table Sizes for OSPF LSAs

A new command in software release 05.0.84 lets you adjust the maximum number of Link-State Advertisements (LSAs) that routing switches with 32MB memory can hold in LSA tables.

The following table lists the types of LSAs for which you can configure table sizes, the default number of entries the tables can hold, and the range of maximum values you can specify. You cannot configure all the LSA tables globally; you must configure them for individual LSA types.

LSA Type	Default Maximum Number of Entries	Range of Values
External	2000	500 - 8000
Network	2000	200 - 2000
Router	2200	200 - 2200
Summary	2000	500 - 8000

To change the maximum number of entries for an LSA type, enter the following command at the OSPF CONFIG level of the CLI:

maximum-number-of-lsa external|network|router|summary <value>

For example, to change the maximum number of summary LSA entries from 2000 to 6000, enter the following commands:

```
HP9300(config-ospf-router)# maximum-number-of-lsa summary 6000
HP9300(config-ospf-router)# write memory
HP9300(config-ospf-router)# exit
HP9300(config)# exit
HP9300# reload
```

You must save the change to the startup-config file and reload or reboot. The change does not take effect until your reload or reboot.

Note

When you allocate high values for the LSA types, you reduce the amount of memory available to other protocols. If you need to run more than one memory-intensive protocol simultaneously, make sure you consider the memory needs of each. For example, if you are planning to run OSPF and BGP4 on the same routing switch, you might not want to allocate the maximum memory capacity for all the LSA types. Doing so could take memory away from BGP4 and thus limit configurable BGP4 table sizes such as the maximum number of routes the BGP route table can hold.

Displaying System Defaults

Software release 05.0.84 contains a new command that displays the defaults for system parameters. To display the default information, enter the following command from any level of the CLI:

show default [values]

If you specify "default" but not the optional "values", the default states for parameters that can either be enabled or disabled are displayed. If you also specify "values", the default values for parameters that take a numeric value are displayed.

Here are some examples of the information displayed by these commands. The first example shows the information displayed by the show default command on an HP 9308M routing switch.

```
HP9300# show default
spanning tree disabled
auto sense port speed      port untagged              port flow control on
no username assigned       no password assigned       boot sys flash primary
system traps enabled       sntp disabled              radius disabled
rip disabled               ospf disabled              bgp disabled
```

when ip routing enabled :

```
ip irdp enabled            ip load-sharing enabled    ip proxy arp enabled
ip rarp enabled            ip bcast forward enabled
dvmrp disabled             pim/dm disabled
vrrp disabled              srp disabled
```

when rip enabled :

```
rip type:v2 only           rip poison rev enabled


ipx disabled               appletalk disabled
```

The following example shows the command output when you use the values option on an HP9308M routing switch.

```
HP9308# show default values
sys log buffers:50         mac age time:300 sec       telnet sessions:5


ip arp age:20 min          bootp relay max hops:4     ip ttl:64 hops
```

```
ip addr per intf:24
when multicast enabled :
igmp group memb.:140 sec   igmp query:60 sec
when ospf enabled :
ospf dead:40 sec           ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100        bgp keep alive:60 sec      bgp hold:180 sec
bgp metric:10              bgp local as:1             bgp cluster id:0
bgp ext. distance:20       bgp int. distance:200      bgp local distance:200
```

System Parameters	Default	Maximum
arp	4000	16000
atalk-route	512	3072
atalk-zone-port	64	255
atalk-zone-sys	255	1024
dvmrp	2048	32000
igmp	255	1024
ip-cache	16000	64000
ip-filter-port	32	256
ip-filter-sys	64	2048
ipx-forward-filter	32	256
ipx-rip-entry	2048	16384
ipx-rip-filter	32	256
ipx-sap-entry	4096	16384
ipx-sap-filter	32	256
l3-vlan	32	1024
ip-qos-session	128	32000
mac	8000	64000
ip-route	10000	200000
ip-static-route	64	1024
vlan	8	4096
mac-filter-port	16	256
mac-filter-sys	32	512

Note

To display defaults information using the Web management interface, select the System link to display the System configuration sheet, then select Parameter. You also can change the parameters from this display.

Listing and Clearing Web Management Sessions

Software release 05.0.84 contains the following new commands to list and clear open Web management sessions:

- **show web-connection** - lists the source IP address and access level ("set" or "get") of each Web management session open on the device
- **clear web-connection** - clears all open Web management sessions

You can enter the show web-connection command from any CLI level. To enter the clear web-connection command you must be at the Privileged EXEC level.

New IP Multicast Commands

The IP multicasting software contains the following new configuration commands:

- **ip multicast-routing** - enables you to configure global IP Multicast routing parameters. Software release 05.0.84 contains the following configurable global IP Multicast parameters:
  - IGMP query interval - As it did prior to release 05.0.84, specifies how often the routing switch queries an interface for group membership. Possible values are 1 - 3600. The default is 60. The CLI command for changing this parameter is **ip igmp group-membership-time <value>**.
  - IGMP group membership time - As it did prior to release 05.0.84, specifies how many seconds an IP Multicast group can remain on a routing switch interface in the absence of a group report. Possible values are 1 - 7200. The default is 60. The CLI command for changing this parameter is **ip igmp query-interval <value>**.
  - IGMP maximum response time (new in software release 05.0.84; see below)
- **ip igmp max-response-time** - Specifies how many seconds the routing switch will wait for an IGMP response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 - 10. The default is 10.

To configure the global IP Multicast parameters listed above, you must enter the ip multicast-routing command. If you do not enter this command, the global parameter changes do not take effect. To ensure that the changes take effect, enter the ip multicast-routing command first, make the global parameter changes, then save the changes to the system-config file.

Trunk Port Configuration Enhancements

Prior to software release 05.0.84, you could not change the port membership in a trunk group unless you deleted the trunk group, then re-added it with a different port membership. For example, to change a port group from ports 1/1 - 1/4 to ports 1/1 - 1/2, you needed to delete the trunk group, then add it again with ports 1/1 and 1/2.

Software release 05.0.84 enables you to change port membership by removing individual ports from the trunk group. To remove a port from a trunk group, enter the following command:

no trunk e <port> [to <port>]

where **<port>** is the port you are removing.

For example, to remove ports 1/3 and 1/4 from the trunk group, enter the following command:

```
HP9300(config)# no trunk e 1/3 to 1/4
```

Note

Make sure you enter the lower port in the range before the "**to**" and the higher port in the range after the "**to**".

As a shortcut, you also can enter just the lower port in the range. The software automatically removes all higher ports in addition to the specified port. For example, to remove ports 1/3 and 1/4, you can enter the following command:

```
HP9300(config)# no trunk e 1/3
```

The rules regarding trunk group membership remain unchanged. Therefore, for trunk group 1/1 - 1/4, the following commands are not valid:

```
HP9300(config)# no trunk e 1/2
```

Or

```
HP9300(config)# no trunk e 1/2 to 1/4
```

These commands are invalid because a trunk group cannot contain only a single port. These commands, if the software allowed them, would result in a trunk group consisting only of port 1/1.

On the HP 9304M and 9308M, trunk groups can contain two ports or four ports but cannot contain only three ports. Therefore, the following command also is invalid for trunk group 1/1 - 1/4:

```
HP9300(config)# no trunk e 1/4
```

This command is invalid because it would result in a trunk group containing three ports, 1/1 - 1/3.

Page-Display Mode Configurable from User EXEC and Privileged EXEC CLI Levels

Page-display mode controls whether the configuration file is presented one "page" (window-full) at a time or all at once when you save or display the file. In software releases earlier than software release 05.0.84, you needed super-user access to change the page-display mode using the enable skip-page-display and no enable skip-page-display commands.

You still can change the page-display mode this way. Software release 05.0.84 add the capability for users at the User EXEC and Privileged EXEC levels of the CLI to change the page-display mode. To change the page-display mode from these CLI levels, use the following commands:

- **page-display** - Enables page-by-page display of the configuration file. This command is equivalent to the no enable skip-page-display command at the global CONFIG level. When you display or save the file, the following line provides you with options to continue the display or to cancel:

  ```
  --More--, next page: Space/Return key, quit: Control-c
  ```

  If you disable the page-display mode, the CLI displays the entire file without interruption. Page-display mode is enabled by default. To disable it, enter the **stop-page-display** command.

- **stop-page-display** - Disables page-display mode. Page-display mode displays the file one page at a time and prompts you to continue or cancel the display. When page-display mode is disabled, if you display or save the configuration file, the CLI displays the entire file without interruption. This command is equivalent to the **no enable skip-page-display** command at the global CONFIG level.

Source IP Addresses Listed in SNMP Authentication Traps

One of the SNMP traps supported by the HP 9304M and 9308M is for SNMP authentication failures. If someone enters an invalid community string when attempting to access the SNMP server on the HP device, the device generates a trap in the device's syslog buffer. (If you have configured the device to use a third-party SyslogD server, the device also sends a log entry to the server.)

Prior to software release 05.0.84, the log messages indicated that an authentication failure occurred but did not list the IP address from which the invalid community string came. In software release 05.0.84, log entries for SNMP authentication failures list the IP address of the device that sent the invalid community string. You do not need to perform any configuration to add this feature. If logging is enabled, the message level includes Informational messages, and the SNMP authentication trap is enabled on the device, then the IP addresses are logged.

Here is an example of a log that contains SNMP authentication traps. In this example, someone attempted to access the HP device three times using invalid SNMP community strings. The unsuccessful attempts indicate either an authorized user who is also a poor typist, or an unauthorized user who is attempting to access the device.

```
HP9300(config)#show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 1 overruns)
Buffer logging: level ACDMEINW, 50 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning

Log Buffer (50 entries):
at 0 days 3 hours 57 minutes 18 seconds, level informational
SNMP Authentication failure, intruder IP: 207.95.6.55
at 0 days 3 hours 57 minutes 14 seconds, level informational
SNMP Authentication failure, intruder IP: 207.95.6.55
at 0 days 3 hours 57 minutes 10 seconds, level informational
SNMP Authentication failure, intruder IP: 207.95.6.55
```

New CLI Command for Displaying Running-Config File

The running-config file contains system configuration information, including changes that you have made during the current management session but have not yet saved to flash memory. Software release 05.0.84 provides a new command for displaying an HP 9304M or 9308M running-config file:

```
show running-config
```

The syntax of the new command makes it more consistent with other display commands such as **show configuration**, which displays the startup-config file. The new command and the older command, **write terminal**, are both listed in the CLI and are both supported.

Configurable MAC Table Size

Prior to release 05.0.84, you could use the following command to configure the sizes of some tables on the HP 9304M and 9308M using the following command:

```
system-max <parameter> <value>
```

These parameters and values for this command are listed under the description of the system-max command on page B-80 in Appendix B, "CLI Reference", in the *HP ProCurve Routing Switch 9304M and 9308M Installation and Configuration Guide.*

Note

You must save the change to the startup-config file and reload or reboot. The change does not take effect until your reload or reboot.

After you enter the system-max mac command, if the system does not have enough memory to implement the increased table size, the system will automatically reduce the amount to 8000. In this case the software also displays an error message listing the changed table size.

- If you want to accept the new table size, save the running-config file and reboot.
- If you want to try the command again with a smaller value, press the Up arrow key to redisplay the command, edit the value, then press Enter.

To return to the default value for the MAC table size, enter the following command. Make sure you save the change to the system-config file and then reboot to implement the change.

no system-max mac <value>

TAB Key Displays Options List

Software release 05.0.84 provides a new method for displaying valid options for a partially completed command string. While both release 05.0.84 and prior software displays an option list when you enter a question mark (?), release 05.0.84 also displays the options when you press the TAB key.

Global-Tag Command Removed from Software

Previous software contains the CLI command **global-tag**. This command enables the ability to configure IEEE 802.1q tagging on ports.

Customer experience indicates that there is no reason to disable the ability to configure tagging. As a result, the **global-tag** command is no longer needed. Tagging is always allowed. To simplify the startup-config file, which sometimes contains this command, the command has been removed from the software in release 05.0.84. For backward compatibility, you still can use older startup-config files that contain the command but the command will no longer appear in startup-config files you save when running software release 05.0.84.

Software Fixes

This section lists the problems that have been fixed in software release 05.0.84.

- **ARP and MAC tables -** When a device attached to an HP 9304M and 9308M routing switch on a Gigabit fiber link went down, the software did not properly detect that the device had gone down and did not clear the MAC and ARP entries for the port connected to the device.
- **Tagging -** In cases where two virtual interfaces contained the same tagged port, Proxy ARP did not send the correct VLAN ID.

- **OSPF -** In prior software, if you cleared the IP route table (using the clear ip route CLI command, for example), the routing switch did not begin learning routes again until you disabled and then re-enabled OSPF. In software release 05.0.84, the software automatically begins learning routes again after you clear the IP route table. You do not need to disable and re-enable OSPF.
- **OSPF (HP 9308M only) -** In some configurations, OSPF did not propagate routes through virtual interfaces.
- **DVMRP -** The Router Alert Option caused an interoperability issue with 3Com® routers. This issue has been fixed and the HP 9304M and 9308M now interoperate with 3Com routers.
- **PIM -** When multiple PIM routers were configured in the same VLAN, the software did not forward multicast packets from one router to the other within the VLAN.
- **IGMP) -** The software did not verify that IGMP reports contained a group address from 224.0.1.0 - 239.255.255.255 (the dynamic range). In software release 05.0.84, the software does not accept a report outside this range.
- **IGMP -** The routing switch did not accept Group Membership reports with TTL great than 1.
- **Telnet -** On rare occasions, an outgoing Telnet session did not properly handle the Telnet Interop Process command, resulting in a crash.
- **User Interface -** When a system date later than February 29, 2000 was configured for testing, the software subtracted a day and thus displayed a date that was one day earlier than the correct date. This problem did not affect the operation of the device and would not have affected the operation of the device even after February 29, 2000; this problem affected only the display of the date. In software release 05.0.84 the date will increment correctly from February 28 to February 29 and from February 29 to March 1, and so on.

HP Part Number: 5969-3509
Edition 1, August 1999
Printed in USA