Release Notes *for* Version 06.6.05 of the HP ProCurve Routing Switch 9304M, 9308M, and 6308M-SX, and Switch 6208M-SX Operating Systems

These release notes describe:

- New operating system enhancements not available in software releases prior to version 06.6.05
- Earlier software operating problems fixed in version 06.6.05

Contents

Software Release 06.6.05

Software version 06.6.05 (or later) and the corresponding product documentation can be downloaded from HP's ProCurve website as described below.

To Download Software Version 06.6.05:

1. Go to HP's ProCurve website at **http://www.hp.com/go/procurve**
2. Click on **Free Software Updates**.
3. Select **Switches download page**.
4. Click on the name of the device for which you are downloading software.

Note: If you are downloading software for a 9304M or 9308M, select the option that matches the type of management module(s) you are using in the device (*with redundant management* or *without redundant management*).

To Download Product Documentation Supporting Software Version 06.6.05:

1. Go to HP's ProCurve website at **http://www.hp.com/go/procurve**
2. Click on **Technical Support**, then **Manuals**.
3. In the resulting display, step through the selections for accessing and viewing the new manuals.
4. On the page listing the manuals, find the new manuals under the heading "**For software version 06.6.00 or greater**".

You will need the Adobe® Acrobat® Reader to view and/or print the manuals. You can also order photocopied, 3-hole punched versions of the new manuals in an 8-1/2 by 11- inch page size and/or the CD-ROM containing PDF files of the manuals. To do so, see the ordering information provided in edition 5 (or later) of the *Read Me First for the HP ProCurve Routing Switches 9304M, 9308M, and 6308M-SX and the HP ProCurve Switch 6208M-SX*. This document is shipped with your HP device and the latest version is also available on the HP ProCurve website. (See steps 1 - 4, above.)

Table 1. Software Download Files for Release 06.6.05

H2R06600.BIN	HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) *With Redundant Management Module(s)*
HPR06600.BIN	HP ProCurve Routing Switch 9304M (J4139A), 9308M (J4138A), and 6308M-SX (J4840A) *Without Redundant Management Module(s)*
HPS06600.BIN	HP ProCurve Switch 6208M-SX (J4841A)

Table 2. Device Compatibility with Software Versions

Device	**Supported Software Versions:**					
	04791	**05084**	**H2R05216.BIN H2R06600.BIN**	**HPR05216.BIN HPR06600.BIN**	**HPS05216.BIN**	**HPS06600.BIN**
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) *With Redundant Management Module(s)*	*No*	*No*	**Yes**	*No*	*No*	*No*
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) *Without Redundant Management Module(s)*	**Yes**	**Yes**	*No*	**Yes**	*No*	*No*
HP ProCurve Routing Switch 6308M-SX (J4840A)	*n/a*	*n/a*	*No*	**Yes**	*No*	*No*
HP ProCurve Switch 6208M-SX (J4841A)	*n/a*	*n/a*	*No*	*No*	**Yes**	**Yes**

Note: The flash image files for these software releases differ depending on the product.

If you have a 9304M or 9308M routing switch that was shipped before version 06.6.05 was available, you may want to download this release from HP's ProCurve website. To do so, see "To Download Software Version 06.6.05", above.

For information on how to update your routing switch software, refer to the chapter titled "Updating Software Images and Configuration Files" in the documentation you received with the device.

Already Using a 9304M, 9308M, 6308M-SX, or 6208M-SX? Here's New Information!

If you received one of the above devices before software release 06.6.05 began shipping, and you are updating the device to release 06.6.05, then you may want to examine the new product manuals that are available beginning with the 06.6.05 release. To view (and freely download) PDF versions of these manuals (chapter-by-chapter files), go to HP's ProCurve website at **http://www.hp.com/go/procurve**, then:

1. Click on **Technical Support**, then **Manual**.
2. In the resulting display, step through the selections for accessing and viewing the new manuals.
3. On the page listing the manuals, find the new manuals under the heading "**For software version 06.6.05 or greater**".

You will need the Adobe® Acrobat® Reader to view and/or print the manuals. You can also order printed versions of either new manual and/or the CD-ROM containing PDF files of both manuals. To do so, see the ordering information provided in edition 5 (or later) of the *Read Me First for the HP ProCurve Routing Switches 9304M, 9308M, and 6308M-SX and the HP ProCurve Switch 6208M-SX*. This document is available on the HP ProCurve website (URL shown above).

Summary of New Features in Release 06.6.05

NOTE: Beginning with software release 06.6.05, the software does not have a default read-write SNMP community. If you use the default community name "private" as the password for Web management access or for read-write access through a network management application, you need to use the CLI to add the read-write community string first.

The following sections list the enhancements in software release 06.6.05.

Most of the 06.6.05 enhancements are described in detail in the Software Release 6.*X* documentation shipped with products containing this version of the software. To get copies of the latest product documentation, see "To Download Product Documentation Supporting Software Version 06.6.05.05" on page 2.

New Hardware Support: Software release 06.6.05 adds support for the new J4842A HP ProCurve 9300 1000Base-T Module (8-port).

Enhancements Added in 06.6.05

Software release 06.6.05 contains the following enhancements.

Layer 3 Enhancements

These enhancements apply only to the 9304M, 9308M, and 6308M-SX routing switches.

- "Null" static IP routes – You can configure a "null" (sometimes called "null0") static route for IP traffic that you want the routing switches to discard. The routing switch drops traffic directed to the specified IP network or host address instead of forwarding it.
- IP static route enhancements:
  - Support for multiple static routes to the same destination – You can configure multiple static routes, with

differing metrics and next-hop gateway addresses, to the same destination.

    - Static routes follow port states – If the HP device's interface associated with a static route is down, the HP software reflects the interface state by making the static route invalid.

- IP load sharing is enabled by default – IP load sharing is enabled by default in software release 06.6.05. The feature is disabled by default in earlier software releases.
- Option to disable forwarding of IP source-routed packets – HP devices forward source-routed IP packets by default but you can disable the forwarding so that the device drops source-routed packets.
- Enhancement to the **show ip interface** command – The output of this command has changed. The command now displays a table summarizing IP status and configuration information for all interfaces.
- Support for default network route – You can configure a default network route for use when the IP route table does not contain any other learned or static default routes.
- Enhancements to **show ip route** command – New options with this command let you restrict the display to static routes only or direct routes only. In addition, you can display longer routes for a specified IP address and network mask.
- Enhancements to **clear ip route** command – You can clear individual routes without clearing the entire table.
- Change to default for forwarding redirected broadcasts – Forwarding of directed IP broadcasts is disabled by default in software release 06.6.05. In earlier software releases the feature is enabled by default.
- Default router ID is the lowest numbered loopback interface – In earlier software releases, the default router ID is the lowest numbered IP address on the routing switch. In software release 06.6.05, the default router ID is the lowest loopback interface on the routing switch.
- Change to IP static route configuration – You do not need to specify an index number when configuring a static IP route.
- Enhanced Border Gateway Protocol version 4 (BGP4) and Open Shortest Path First (OSPF) route redistribution options – When you configure a routing switch to perform redistribution of BGP4 or OSPF routes, you can specify the types of routes (connected, RIP, OSPF, or static) that you want to redistribute into BGP4 or OSPF.
- BGP4 enhancements:
    - BGP4 Autonomous System (AS) confederations – You can logically subdivide an AS into a confederation of multiple ASs, thus allowing the routing switch to be used in large ASs and simplifying the Interior Border Gateway Protocol (IBGP) mesh among the routing switches within the AS. The HP implementation of this feature is based on RFC 1965.
    - BGP4 load sharing – You can configure the routing switch to balance IP traffic to a given destination across up to eight equal IBGP or EBGP paths (all the paths must either be IBGP or EBGP.)
    - Changes to BGP4 commands and default values – The CLI syntax and the defaults for some BGP4 commands and parameters have changed.
    - BGP4 start-failure event and neighbor state change events – The routing switch generates an SNMP trap and Syslog message if the BGP4 process on a routing switch fails to start properly. In addition, the routing switch generates Syslog message and SNMP traps when a neighbor's state changes, including when the neighbor comes up or goes down.
    - Route flap dampening
    - Advanced BGP4 route filtering using ACLs and prefix lists
    - Route map enhancements for matching on AS-paths, communities, destination networks, and next-hop routers
    - Neighbor configuration enhancements
    - New commands and Web management options to display AS-path ACLs, community lists, IP prefix lists, and route filter lists
    - Help messages and performance enhancements for regular expressions

    - "Longer" option with **show ip bgp** command
- OSPF enhancements
    - Default route origination (default-information-originate) is disabled by default
    - Default route advertisement regardless of redistribution
    - No-summary option for stub areas – You can disable advertisement of summary Link State Advertisements (LSAs) into stub areas.
    - Route summarization – You can summarize routes redistributed into OSPF to reduce the number of External LSAs.
    - Default-information originate option – The routing switch generates a default route into the OSPF domain by default.
    - Configurable administrative distances for OSPF route types – You can change the default administrative distances for OSPF inter-area, intra-area, and external routes. The default administrative distance for all these routes is still 110.
    - Blocking of LSA floods – You can block flooding of outbound LSAs on individual interfaces.
    - Configurable Shortest Path First (SPF) calculation timers – You can change the SPF delay and SPF hold time timers.
    - OSPF Not So Stubby Areas (NSSA) – You can configure NSSAs per RFC 1587. An NSSA allows the area to contain Autonomous System Border Routers (ASBRs) that import External Link State Advertisements (LSAs), while at the same time prohibiting flooding of External LSAs into the NSSA.
    - OSPF passive interface option – You can configure an OSPF interface to not advertise route information.
    - The **normal** parameter is no longer valid with the OSPF **area** command – The software assumes that an OSPF area is normal unless you specify otherwise. Thus, the unnecessary **normal** parameter has been removed.
    - OSPF load sharing is enabled by default – The default state of OSPF load sharing has been changed from disabled to enabled. The default maximum number of paths is still four.
- IP Multicast Protocol-Independent Multicast (PIM) enhancements:
    - Configure static PIM routes, query another PIM router for its PIM configuration information, and trace a PIM route to its source.
    - Adds support for PIM sparse, which provides multicast service for widely dispersed multicast environments.
- IPX Server Advertisement Protocol (SAP) access lists – You can configure up to 32 access lists to control the routing switch's responses to SAP requests on specific interfaces.
- IPX Get Nearest Server (GNS) enhancements – You can configure GNS filters to filter the servers in GNS replies, use round robin to rotate among a list of servers when sending GNS replies, and disable an interface from sending GNS replies.
- VRRP enhancements:
    - You can force a failover of a Master router to a Backup router. In addition, the **show ip vrrp** command has been enhanced to list the IP addresses of the Backup routers.
    - Changes to keepalive parameter for VRRP backup routers – This release changes the default keepalive state to disabled, which means backup routers do not send keepalive messages by default. In addition, if you want to use the keepalive messages, you can configure the message interval.
- AppleTalk ARP age – You can now set this parameter globally using the CLI. You no longer need to set the parameter on an individual interface basis.
- Enhancements to virtual interfaces – You can now administratively disable virtual interfaces.

Layer 2 Enhancements

Software release 06.6.05 contains enhancements to the Spanning Tree Protocol (STP). Unless otherwise noted, these enhancements apply both to the 6208M-SX switch and to the 9304M, 9308M, and 6308M-SX routing switches.

- Fast Port Span – By default, devices running software release 06.6.05 perform Spanning Tree Protocol (STP) convergence in four seconds instead of 30 or more seconds for certain ports connected to end stations.
- Fast Uplink – This feature enhances STP by allowing an HP device with redundant uplinks to quickly resume forwarding, in just four seconds. This feature is similar to Fast Port Span but applies to certain inter-switch links on HP devices, instead of HP links to end stations.
- Single-instance STP – You can configure an HP device to run a single spanning tree, even if you have already configured multiple port-based VLANs on the device.

System Level Enhancements

Software release 06.6.05 contains the following system-level enhancements. Unless otherwise noted, these enhancements apply both to switches and to routing switches.

- Enhanced security against Smurf and TCP SYN Denial of Service (DoS) attacks
- Interface-based static routes
- Support for IP zero-based IP sub-net broadcasts
- Change of default IP proxy ARP state from enabled to disabled
- Access Control Lists (ACLs) – You can configure ACLs and apply them to the incoming or outgoing traffic on individual interfaces to filter packets based on IP and TCP/UDP information.
- Quality of Service (QoS) – You can configure QoS profiles that describe bandwidth and prioritization settings, and QoS groups that describe the attributes of incoming traffic that are affected by the QoS profiles. These new QoS features expand the QoS features in previous releases.
- Command Line Interface (CLI) changes – The syntax for some CLI commands has been changed to more closely resemble similar command syntax on third-party devices.
- CLI context-sensitive help – You can display a list of all the commands and options at a given CLI level that begin with a specific character or string of characters. In addition, the CLI contains descriptions of the commands.
- Web management interface enhancements – The Web management interface contains navigation and display enhancements to make the interface easier to use.
- New command for disabling Web management access authentication – You can disable authentication for Web management access, allowing you to access the interface without specifying “get’ or “set” and a community string.
- Syslog message enhancements – The format of Syslog messages is changed to enhance readability while allowing more messages to be displayed on the screen at one time.
- SNMPv2c support – The SNMP agent on HP devices now supports SNMPv2, including GetBulk requests.
- Routing switches only – The maximum IP ARP table size has been increased to hold 64,000 entries. The default maximum size of the table is still 8,000 entries.
- Enhancement to **reload** command – When you enter this command, the device asks you to verify whether you want to reload the software. In addition, if you have made configuration changes that are not saved in the system-config file, the software gives you an opportunity to save the changes first before reloading.
- Enhancement to interface configuration displays – The CLI command and Web management panel for displaying interface configuration information now provides additional information, including packet counters and input and output rate statistics. The command and option also can display information about loopback interfaces and virtual ethernet (VE) interfaces.
- Enhancement to **show cpu** command – This command now can report the percentage of CPU usage over a

specified interval.

- Configurable CLI greeting banners
- Increase in number of SyslogD servers supported from two to six
- Support for configuring a loopback interface as an SNMP trap source, in addition to support in previous software releases for configuring an Ethernet port or virtual interface as the trap source

---

NOTE: This enhancement applies only to routing switches.

- Configurable source IP address for Telnet and TACACS/TACACS+ packets originated from the routing switch

---

NOTE: This enhancement applies only to routing switches.

- Additional CLI enhancements:
  - When you display commands in the history buffer (by pressing Ctrl-N, Ctrl-P, or the up and down arrows), the index number no longer appears to the left of the command.
  - The Ctrl-Z key takes you from any CONFIG level of the CLI to the Privileged EXEC level. If you are at the Privileged EXEC level, Ctrl-Z takes you to the User EXEC level.
  - If you enter an invalid command followed by ?, the CLI displays the “Unrecognized command“ message to indicate that the command does not exist.
- Encryption of the display of SNMP community strings – SNMP community strings are displayed in encrypted format in the CLI by default. In addition, users with read-only access cannot view the SNMP community strings in the Web management interface. If you want SNMP community strings to be displayed in the clear, you can disable encryption on an individual community string basis.
- IP ACL enhancements:
  - Fully supported on chassis routing switches
  - Support for IP Quality-of-Service (QoS) options such as precedence and Type of Service (TOS)
  - Support for ACL names
  - Supported for BGP4 filtering and redistribution
- Encrypted display of SNMP community strings – Display of SNMP community strings is encrypted by default beginning with software release 06.6.05. However, if you prefer to display the strings in the clear, the software allows you to specify this.
- Option to kill Telnet sessions – An administrator with super-user access can kill active Telnet CLI management sessions.
- CLI serial session timeout – You can configure a device to time out an idle CLI session on the serial port.
- Route-only option available on individual interface basis – you can disable all Layer 2 switching on an individual interface basis. In previous releases, you can disable Layer 2 switching only on a global basis.
- IP ping sent to the IP broadcast address – When you ping to the IP broadcast address, the first four responses are displayed by default. In previous releases, if you address the ping to the IP broadcast address, the device lists the responses from all systems that reply to the ping.
- Enhancements to the Traceroute function – Traceroute requests display all responses to a given TTL. In addition, if there are multiple equal-cost routes to the destination, the HP device displays up to three responses by default.
- Syslog enhancements:
  - You can dynamically change the size of the Syslog buffer; this no longer requires you to reload the software.
  - You can enable logging on a configured ACL or filter that supports logging by re-entering the ACL or filter

configuration command and adding **log** to the end.

  - Power supply and temperature sensor log entries reside in a separate, permanent buffer.
  - Syslog messages contain more interface information.
  - Traps that occur while an interface is coming up are saved and logged when the interface comes up.
  - Changes saved to the startup-config file are logged.
  - New log messages log RIP filter activity.
- CLI enhancements:
  - You can clear statistics on individual modules or ports.
  - The command prompt at the Interface configuration level shows the port speed for Ethernet ports.
  - Commands such as **ip route** that accept an IP network mask as a parameter accept Classless Interdomain Routing (CIDR) notation (for example, 209.157.22.110/24 instead of 209.157.22.110 255.255.255.0). In previous releases, only some commands supported the CIDR notation.
  - The **show running-config** command display SNMP community strings and passwords when the command is entered from the Privileged EXEC mode, but does not display them when entered from the User EXEC mode.
  - The **show interfaces** command displays detailed information for all interfaces. A new command, **show interfaces brief**, shows only the Layer 2 information for the interfaces.
- Chassis software performance tuning to enhance forwarding performance for certain routing protocols.
- The maximum number of Virtual Router Redundancy Protocol (VRRP) virtual router IDs (VRIDs) is increased from 4 to 12.

Software Image Files

To run software release 06.6.05, you need the boot and flash images listed in the following table.

Product	Flash Image
9308M 9304M	Redundant Management modules: • H2R06605.bin (routing switch code) Other management modules: • HPR06605.bin (routing switch code)
6308M-SX	HPR06605.bin (routing switch code)
6208M-SX	HPS06605.bin (switch code)

NOTE: To upgrade redundant management modules in a 9304M or 9308M chassis, use the instructions in the "Using Redundant Management Modules" chapter of *Book 1: Installation and Getting Started Guide*. To get the latest version of this guide, see "To Download Product Documentation Supporting Software Version 06.6.05" on page 2.

NOTE: If you are adding a 1000Base-T module to a 9304M or 9308M chassis you must upgrade to flash image H2R06600 (if your are using redundant management modules) or HPR06600 (if you are *not using* redundant management modules).

Logging On

You can access an HP device 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: The serial cable is shipped with your switch or routing switch. If you prefer to build your own cable, see the pinout information in the Installation chapter of *Book 1: Installation and Getting Started Guide*. To get the latest version of this guide, see "To Download Product Documentation Supporting Software Version 06.6.05" on page 2.

Accessing a Routing Switch CLI

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

NOTE: This procedure applies only to routing switches. Use the next procedure for switches.

1. Enter the **enable** command at the User EXEC level prompt (for example, `HP9300>`). Then press Enter. (If you are prompted for a password, enter your enable password.)
2. Enter the **erase startup-config** command at the Privileged EXEC level prompt (for example, `HP9300#`), then press Enter. This command erases the factory test configuration if still present.

CAUTION: Use this step only for new systems. If you enter this command on a system you have already configured, the command erases the configuration. If you accidentally do erase the configuration on a configured system, enter the **write memory** command to save the running configuration to the startup-config file.

3. Enter the **configure terminal** command at the Privileged EXEC level prompt. Then press Enter.
4. Enter the following command to access the Interface CONFIG level for the interface:
   int e <*portnum*>, where <*portnum*> is the port number of the interface.
5. Enter the following command to configure the IP address:
   ip address <*ip-addr*> <*ip-mask*>

   NOTE: You also can enter **ip address** <*mask-bits*>.
6. Enter the **write memory** command to save the configuration information to the HP device's flash memory.

NOTE: You will be able to access the routing switch only through this physical interface.

Accessing a Switch CLI

To access a switch's CLI through a Telnet connection, assign an IP address to the switch as follows.

NOTE: This procedure applies only to switches. Use the procedure above for routing switches.

1. Enter the **enable** command at the User EXEC level prompt (for example, `HP9300>`). Then press Enter. (If you are prompted for a password, enter your enable password.)
2. Enter the **erase startup-config** command at the Privileged EXEC level prompt (for example, `HP9300#`), then press Enter. This command erases the factory test configuration if still present.

CAUTION: Use this step only for new systems. If you enter this command on a system you have already configured, the command erases the configuration. If you accidentally do erase the configuration on a configured system, enter the **write memory** command to save the running configuration to the startup-config file.

---

3. Enter the **configure terminal** command at the Privileged EXEC level prompt. Then press Enter.
4. Enter the **ip address** <*ip-addr*> <*ip-mask*> command at the device's global CONFIG level prompt (for example, `HP9300(config)#`). Then press Enter.
5. To set a default gateway address (optional), enter the **ip default gateway** <*ip-addr*> command, then press Enter.
6. Enter the **write memory** command to save the configuration information to the HP device's flash memory.

Accessing the Web Management Interface

To access the Web management interface through a Web browser, enter the HP 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 a valid read-write community string in the Password field. (If this does not work, you have not yet configured the read-write community string using the CLI. Beginning with software release 05.2.16, there is no default read-write community string.)
- 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 device.)

Recovering from a Lost Password

By default, the CLI does not require passwords. However, if someone has configured a password for the HP 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.) This command causes the device to bypass the system password check.
5. Enter **boot system flash primary** at the prompt.
6. After the console prompt reappears, assign a new password.

Enhancements in Release 06.6.05

Most of the enhancements in release 06.06.05 are described in the updated product documentation. If you need to acquire new documentation, see "To Download Product Documentation Supporting Software Version 06.6.05" on page 2. This section covers enhancements not included in the updated product documentation.

Protection Against Denial of Service (DoS) Attacks

In a Denial of Service (DoS) attack, a router is flooded with useless packets, hindering normal operation. Release 06.6.05 includes measures for defending against two types of DoS attacks: Smurf attacks and TCP SYN attacks.

Protecting Against Smurf Attacks

A *Smurf attack* is a kind of DoS attack where an attacker causes a victim to be flooded with ICMP (Ping) replies sent from another network. Figure 1 illustrates how a Smurf attack works.

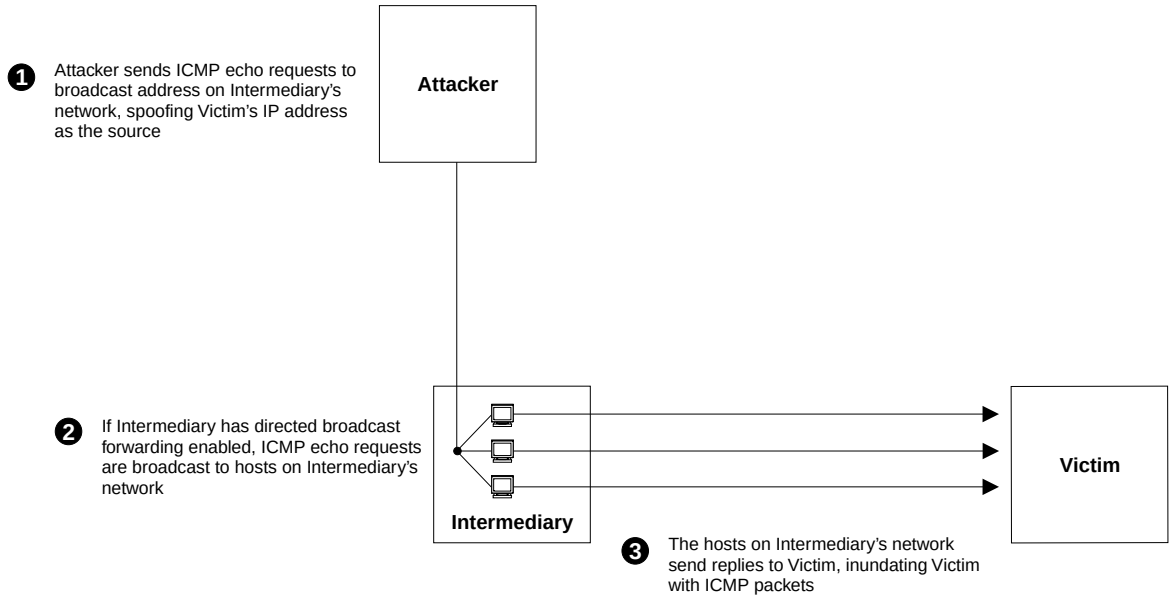

Figure 1 How a Smurf attack floods a victim with ICMP replies

The attacker sends an ICMP echo request packet to the broadcast address of an intermediary network. The ICMP echo request packet contains the spoofed address of a victim network as its source. When the ICMP echo request reaches the intermediary network, it is converted to a Layer 2 broadcast and sent to the hosts on the intermediary network. The hosts on the intermediary network then send ICMP replies to the victim network.

For each ICMP echo request packet sent by the attacker, a number of ICMP replies equal to the number of hosts on the intermediary network are sent to the victim. If the attacker generates a large volume of ICMP echo request packets, and the intermediary network contains a large number of hosts, the victim can be overwhelmed with ICMP replies.

Avoiding Being an Intermediary in a Smurf Attack

A Smurf attack relies on the intermediary to broadcast ICMP echo request packets to hosts on a target sub-net. When the ICMP echo request packet arrives at the target sub-net, it is converted to a Layer 2 broadcast and sent to the connected hosts. This conversion takes place only when directed broadcast forwarding is enabled on the device.

To avoid being an intermediary in a Smurf attack, make sure forwarding of directed broadcasts is disabled on the HP device. Starting with release 06.6.05, directed broadcast forwarding is disabled by default. In releases prior to 06.6.05, directed broadcast forwarding is enabled by default. To disable directed broadcast forwarding, do one of the following:

USING THE CLI

```
HP9300(config)# no ip directed-broadcast
```

Syntax: no ip directed-broadcast

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to IP to display the list of IP configuration options.
4. Select the General link to display the IP configuration panel.
5. Select Disable next to Directed Broadcast Forward.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Avoiding Being a Victim in a Smurf Attack

Starting with Release 06.6.05, you can configure the HP device to drop ICMP packets when excessive numbers are encountered, as is the case when the device is the victim of a Smurf attack. You can set threshold values for ICMP packets targeted at the router and drop them when the thresholds are exceeded.

For example:

```
HP9300(config)# ip icmp burst-normal 5000 burst-max 10000 lockup 300
```

Syntax: ip icmp burst-normal <*value*> burst-max <*value*> lockup <*seconds*>

The number of incoming ICMP packets per second are measured and compared to the threshold values as follows:

- If the number of ICMP packets exceeds the **burst-normal** value, the excess ICMP packets are dropped.
- If the number of ICMP packets exceeds the **burst-max** value, *all* ICMP packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

In the example above, if the number of ICMP packets received per second exceeds 5,000, the excess packets are dropped. If the number of ICMP packets received per second exceeds 10,000, the device drops all ICMP packets for the next 300 seconds (five minutes).

When incoming ICMP packets exceed the **burst-max** value, the following message is logged:

<*date*> <*time*>:N:Local ICMP exceeds <*burst-max*> burst packets, stopping for <*lockup*> seconds!!

Protecting Against TCP SYN Attacks

TCP SYN attacks exploit the process of how TCP connections are established in order to disrupt normal traffic flow. When a TCP connection starts, the connecting host first sends a TCP SYN packet to the destination host. The destination host responds with a SYN ACK packet, and the connecting host sends back an ACK packet. This process, known as a "TCP three-way handshake", establishes the TCP connection.

While waiting for the connecting host to send an ACK packet, the destination host keeps track of the as-yet incomplete TCP connection in a connection queue. When the ACK packet is received, information about the connection is removed from the connection queue. Usually there is not much time between the destination host sending a SYN ACK packet and the source host sending an ACK packet, so the connection queue clears quickly.

In a TCP SYN attack, an attacker floods a host with TCP SYN packets that have random source IP addresses. For each of these TCP SYN packets, the destination host responds with a SYN ACK packet and adds information to the connection queue. However, since the source host does not exist, no ACK packet is sent back to the destination host, and an entry remains in the connection queue until it ages out (after around a minute). If the attacker sends enough TCP SYN packets, the connection queue can fill up, and service can be denied to legitimate TCP connections.

To protect against TCP SYN attacks, you can configure the HP device to drop TCP SYN packets when excessive numbers are encountered. You can set threshold values for TCP SYN packets targeted at the router and drop them when the thresholds are exceeded.

For example:

```
HP9300(config)# ip tcp burst-normal 10 burst-max 100 lockup 300
```

Syntax: ip tcp burst-normal <*value*> burst-max <*value*> lockup <*seconds*>

The number of incoming TCP SYN packets per second are measured and compared to the threshold values as follows:

- If the number of TCP SYN packets exceeds the **burst-normal** value, the excess TCP SYN packets are dropped.
- If the number of TCP SYN packets exceeds the **burst-max** value, *all* TCP SYN packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

In the example above, if the number of TCP SYN packets received per second exceeds 10, the excess packets are dropped. If the number of TCP SYN packets received per second exceeds 100, the device drops all TCP SYN packets for the next 300 seconds (five minutes).

When incoming TCP SYN packets exceed the **burst-max** value, the following message is logged:

<*date*> <*time*>:N:Local TCP exceeds <*burst-max*> burst packets, stopping for <*lockup*> seconds!!

Interface-Based Static Routes

Software release 06.6.05 allows you to configure IP static routes by specifying an Ethernet port or a virtual interface instead of a next-hop (gateway) address. For such routes, the routing switch forwards packets destined for the static route's destination network to the specified interface. Conceptually, this feature makes the destination network like a directly connected network, associated with a specific routing switch interface.

NOTE: The port or virtual interface you use for the static route must have at least one IP address configured on it. The address does not need to be in the same sub-net as the destination network.

To configure a static route with a port or interface instead of a next-hop address, use either of the following methods.

USING THE CLI

To configure a static IP route with an Ethernet port instead of a next-hop address, enter a command such as the following.

```
HP9300(config)# ip route 192.128.2.69 255.255.255.0 ethernet 4/1
```

The command in the example above configures a static IP route for destination network 192.128.2.69/24. Since an Ethernet port is specified instead of a gateway IP address as the next hop, the routing switch always forwards traffic for the 192.128.2.69/24 network to port 4/1.

Syntax: [no] ip route <*ip-addr*> <*ip-mask*> ethernet <*portnum*> | ve <*num*>
[<*metric*>] [distance <*num*>]

or

Syntax: [*no*] ip route <*ip-addr*>/<*mask-bits*> ethernet <*portnum*> | ve <*num*> [<*metric*>] [distance <*num*>]

The <*num*> parameter is a virtual interface number. If you instead specify an Ethernet port, the <*portnum*> is the port's number (including the slot number, if you are configuring a 9304M or 9308M).

The command in the following example configures an IP static route that uses virtual interface 3 as its next-hop.

```
HP9300(config)# ip route 192.128.2.71 255.255.255.0 ve 3
```

The command in the following example configures an IP static route that uses port 2/2 as its next-hop.

```
HP9300(config)# ip route 192.128.2.73 255.255.255.0 eth 2/2
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.
4. Click on the General link to display the IP configuration panel.
5. Click the Static Route link.
   - If the device does not have any IP static routes, the Static Route configuration panel is displayed.
   - If a static route is already configured and you are adding a new route, click on the Add Static Route link to display the Static Route configuration panel.
   - If you are modifying an existing static route, click on the Modify button to the right of the row describing the static route to display the Static Route configuration panel.
6. Enter the network address for the route in the Network field.
7. Enter the network mask in the Mask field.
8. Select the next-hop type. You can select one of the following:
   - Address – The next-hop is the IP address of a gateway router.
   - Interface – The next hop is a port, loopback interface, or virtual interface on the routing switch.
9. Enter the next-hop IP address (if you selected the Address method) or select the interface (if you selected the Interface method).
   - Address – Enter the IP address of the next-hop gateway in the Next Hop (by Address) field.
   - Interface – Select the port, loopback interface, or virtual interface from the Next Hop (by Interface) field's pulldown menu(s). Loopback interfaces and virtual interfaces are listed in the Port pulldown menu, not in the Slot pulldown menu. To select a loopback interface or a virtual interface on a chassis device, ignore the Slot pulldown menu and select the interface from the Port pulldown menu.
10. Optionally change the metric by editing the value in the Metric field. You can specify a number from 1 – 16. The default is 1.

    NOTE: If you specify 16, RIP considers the metric to be infinite and thus also considers the route to be unreachable.

    ---

11. Optionally change the administrative distance by editing the value in the Distance field. When comparing otherwise equal routes to a destination, the routing switch prefers lower administrative distances over higher ones, so make sure you use a low value for your default route. The default is 1.
12. Click the Add button to save the change to the device's running-config file.

13. Repeat steps 6 – 10 for each static route to the same destination.
14. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Support for Zero-Based IP Sub-Net Broadcasts

By default, the routing switch treats IP packets with all ones in the host portion of the address as IP broadcast packets. For example, the routing switch treats IP packets with 209.157.22.255/24 as the destination IP address as IP broadcast packets and forwards the packets to all IP hosts within the 209.157.22.x sub-net (except the host that sent the broadcast packet to the routing switch).

Most IP hosts are configured to receive IP sub-net broadcast packets with all ones in the host portion of the address. However, some older IP hosts instead expect IP sub-net broadcast packets that have all zeros instead of all ones in the host portion of the address. To accommodate this type of host, you can enable the routing switch to treat IP packets with all zeros in the host portion of the destination IP address as broadcast packets.

NOTE: When you enable the routing switch for zero-based sub-net broadcasts, the routing switch still treats IP packets with all ones the host portion as IP sub-net broadcasts too. Thus, the routing switch can be configured to support all ones only (the default) or all ones ***and*** all zeroes.

NOTE: This feature applies only to IP sub-net broadcasts, not to local network broadcasts. The local network broadcast address is still expected to be all ones.

To enable the routing switch for zero-based IP broadcasts, use either of the following methods.

USING THE CLI

To enable the routing switch for zero-based IP sub-net broadcasts in addition to ones-based IP sub-net broadcasts, enter the following command.

```
HP9300(config)# ip broadcast-zero
```

Syntax: [no] ip broadcast-zero

USING THE WEB MANAGEMENT INTERFACE

You cannot enable zero-based IP sub-net broadcasting using the Web management interface.

IP Proxy ARP Disabled By Default

Software release 06.6.05 changes the default state of IP proxy ARP from enabled to disabled. When you boot your routing switch with software release 06.6.05, IP proxy ARP is disabled by default instead of enabled. If you want re-enable IP proxy ARP, use either of the following methods.

USING THE CLI

To enable IP proxy ARP, enter the following command:

```
HP9300(config)# ip proxy-arp
```

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

```
HP9300(config)# no ip proxy-arp
```

Syntax: [no] ip proxy-arp

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.

4. Click on the General link to display the IP configuration panel.
5. Select the Enable or Disable radio button next to Proxy ARP.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

OSPF Enhancements

Software release 06.6.05 contains the following OSPF default route enhancements:

- Default route origination (default-information-originate) is disabled by default.
- Default route advertisement. In this release, the routing switch will advertise the default route into OSPF even if redistribution is not enabled, and even if the default route is learned through an IBGP neighbor.

Default Route Origination Is Disabled By Default

When you boot your routing switch with software release 06.6.05, OSPF default route origination is disabled by default instead of enabled. If you want re-enable OSPF default route origination, use the following CLI method.

USING THE CLI

To enable default information originate, enter the following command:

```
HP9300(config-ospf-router)# default-information-originate
```

To again disable the feature, enter the following command:

```
HP9300(config-ospf-router)# no default-information-originate
```

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

USING THE WEB MANAGEMENT INTERFACE

You cannot configure OSPF default information origination using the Web management interface.

Automatic Advertisement of OSPF Default Routes

In software release 06.6.05, the routing switch advertises the default route into OSPF even if redistribution is not enabled, and even if the default route is learned through an IBGP neighbor.

NOTE: IBGP routes (including the default route) are not redistributed into OSPF by OSPF redistribution (for example, by the OSPF **redistribute** command).

Configurable CLI Banners

HP devices can be configured to display a greeting message on users' terminals when they enter the Privileged EXEC CLI level or access the device through Telnet. In addition, the HP device can display a message on the Console when an incoming Telnet CLI session is detected.

Setting a Message of the Day Banner

You can configure the HP device to display a message on a user's terminal when he or she establishes a Telnet CLI session. For example, to display the message "Welcome to the HP 9308M!" when a Telnet CLI session is established:

```
HP9300(config)# banner motd $ (Press Return)
Enter TEXT message, End with the character '$'.
Welcome to the HP 9308M! $
```

A delimiting character is established on the first line of the **banner motd** command. You begin and end the message with this delimiting character. The delimiting character can be any character except " (double-quotation mark) and cannot appear in the banner text. In this example, the delimiting character is $ (dollar

sign). The text in between the dollar signs is the contents of the banner. The banner text can be up to 2048 characters long and can consist of multiple lines. To remove the banner, enter the **no banner motd** command.

Syntax: [no] banner <*delimiting-character*> | [motd <*delimiting-character*>]

NOTE: The **banner** <*delimiting-character*> command is equivalent to the **banner motd** <*delimiting-character*> command.

Setting a Privileged EXEC CLI Level Banner

You can configure the HP device to display a message when a user enters the Privileged EXEC CLI level. For example:

```
HP9300(config)# banner exec_mode # (Press Return)
Enter TEXT message, End with the character '#'.
You are entering Privileged EXEC level #
```

As with the **banner motd** command, you begin and end the message with a delimiting character; in this example, the delimiting character is # (pound sign). To remove the banner, enter the **no banner exec_mode** command.

Syntax: [no] banner exec_mode <*delimiting-character*>

Displaying a Message on the Console When an Incoming Telnet Session Is Detected

You can configure the HP device to display a message on the Console when a user establishes a Telnet session. This message indicates where the user is connecting from and displays a configurable text message.

For example:

```
HP9300(config)# banner incoming $ (Press Return)
Enter TEXT message, End with the character '$'.
Incoming Telnet Session!! $
```

When a user connects to the CLI using Telnet, the following message appears on the Console:

```
Telnet from 209.157.22.63
Incoming Telnet Session!!
```

Syntax: [no] banner incoming <*delimiting-character*>

To remove the banner, enter the **no banner incoming** command.

Increased SyslogD Server Support

Beginning with the 06.6.05 release, six SyslogD servers are supported. Previous versions of the software supported two SyslogD servers. As with previous versions, when multiple SyslogD servers are selected, all the servers receive all Syslog messages sent from the device. The syntax for setting up SyslogD servers remains unchanged.

Syntax: logging <*ip-addr*> | <*server-name*>

Enhanced Traceroute Facility

In past software releases, the **traceroute** command displayed route information only after three attempts to contact each router in the route. In this release, information about a route is displayed after each attempt. The information displayed and the **traceroute** syntax and options remain unchanged. For information on the **traceroute** command, see *Book 2: Advanced Configuration and Management Guide* and *Book 3: Command Line Interface Reference*.

Support for Configuring Loopback Interfaces as SNMP Trap Sources

Previous software releases allow you to configure the routing switch to use the first IP address configured on a specific Ethernet port or virtual interface (VE) as the IP address the HP device uses as the source for all SNMP traps sent by the device.

Software release 06.6.05 extends this support by allowing you to specify the first IP address configured on a loopback interface as the SNMP trap source.

Identifying a single source IP address for SNMP traps provides the following benefits:

- If your trap receiver is configured to accept traps only from specific links or IP addresses, you can use this feature to simplify configuration of the trap receiver by configuring the HP device to always send the traps from the same link or source address.
- If you specify a loopback interface as the single source for SNMP traps, SNMP trap receivers can receive traps regardless of the states of individual links. Thus, if a link to the trap receiver becomes unavailable but the receiver can be reached through another link, the receiver still receives the trap, and the trap still has the source IP address of the loopback interface.

NOTE: When you designate an interface as the SNMP trap source for an HP device, the software uses the first IP address configured on the interface as the source IP address for the traps. The first IP address refers to when the address was configured, not to its numeric sequence relative to other IP addresses configured on the interface. Thus, the first IP address is not the numerically lowest address, but is instead the IP address configured one the interface before any other IP addresses were configured on that interface.

To specify an SNMP trap source, use either of the following methods.

USING THE CLI

To specify the first IP address configured on a loopback interface as the device's SNMP trap source, enter commands such as a the following:

```
HP9300(config)# int loopback 1
HP9300(config-lbif-1)# ip address 10.0.0.1/24
HP9300(config-lbif-1)# exit
HP9300(config)# snmp trap-source loopback 1
```

The commands in this example configure loopback interface 1, assign IP address 10.00.1 to the loopback interface, then designate the interface as the SNMP trap source for this routing switch. Regardless of the port the HP device uses to send traps to the receiver, the traps always arrive from the same source IP address.

Syntax: interface loopback <*num*>

The <*num*> value can be from 1 – 8 on the 9304M, 9308M, and 6308M-SX, and from 1 – 4 on the 6208M-SX.

Syntax: snmp trap-source loopback <*num*> | ethernet <*portnum*> | ve <*num*>

The <*num*> parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the <*portnum*> is the port's number (including the slot number, if you are configuring a 9304M or 9308M).

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a single SNMP trap source using the Web management interface.

Support for Configuring an Interface as the Source for All Telnet or TACACS/TACACS+ Packets

Software release 06.6.05 enables you to designate the first IP address configured an Ethernet port, loopback interface, or virtual interface as the source IP address for all Telnet or TACACS/TACACS+ packets from the routing switch.

Identifying a single source IP address for Telnet or TACACS/TACACS+ packets provides the following benefits:

- If your Telnet client or TACACS/TACACS+ server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the Telnet client or TACACS/TACACS+ server by configuring the HP device to always send the Telnet or TACACS/TACACS+ packets from the same link or source address.
- If you specify a loopback interface as the single source for Telnet or TACACS/TACACS+ packets, Telnet clients or TACACS/TACACS+ servers can receive the packets regardless of the states of individual links. Thus, if a link to the Telnet client or TACACS/TACACS+ server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet and for TACACS/TACACS+. You can configure a source interface for one or both types of packets.

NOTE: The software uses the first IP address configured on the interface as the source IP address for the traps. The first IP address refers to when the address was configured, not to its numeric sequence relative to other IP addresses configured on the interface.

NOTE: When you specify a single Telnet source, you can use only that source address to establish Telnet management sessions with the HP device.

Configuring an Interface as the Source for All Telnet Packets

To specify an Ethernet port or a loopback or virtual interface as the source for all Telnet packets from the device, use either of the following methods. The software uses the first IP address configured on the port or interface as the source IP address for Telnet packets originated by the device.

NOTE: When you specify a single Telnet source, you can use only that source address to establish Telnet management sessions with the HP device.

USING THE CLI

To specify the first IP address configured on a loopback interface as the device's source for all Telnet packets, enter commands such as a the following:

```
HP9300(config)# int loopback 2
HP9300(config-lbif-2)# ip address 10.0.0.2/24
HP9300(config-lbif-2)# exit
HP9300(config)# ip telnet source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all Telnet packets from the routing switch.

Syntax: ip telnet source-interface ethernet <*portnum*> | loopback <*num*> | ve <*num*>

The <*num*> parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the <*portnum*> is the port's number (including the slot number, if you are configuring a 9304M or 9308M).

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the routing switch.

```
HP9300(config)# interface ethernet 1/4
HP9300(config-if-1/4)# ip address 209.157.22.110/24
HP9300(config-if-1/4)# exit
HP9300(config)# ip telnet source-interface ethernet 1/4
```

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a single Telnet source using the Web management interface.

Configuring an Interface as the Source for All TACACS/TACACS+ Packets

To specify an Ethernet port or a loopback or virtual interface as the source for all TACACS/TACACS+ packets from the device, use either of the following methods. The software uses the first IP address configured on the port or interface as the source IP address for TACACS/TACACS+ packets originated by the device.

USING THE CLI

To specify the first IP address configured on a virtual interface as the device's source for all TACACS/TACACS+ packets, enter commands such as a the following:

```
HP9300(config)# int ve 1
HP9300(config-vif-1)# ip address 10.0.0.3/24
HP9300(config-vif-1)# exit
HP9300(config)# ip tacacs source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all TACACS/TACACS+ packets from the routing switch.

Syntax: ip tacacs source-interface ethernet <*portnum*> | loopback <*num*> | ve <*num*>

The <*num*> parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the <*portnum*> is the port's number (including the slot number, if you are configuring a 9304M or 9308M).

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a single TACACS/TACACS+ source using the Web management interface.

Additional CLI Enhancements

- When you display commands in the history buffer (by pressing Ctrl-N, Ctrl-P, or the up and down arrows), the index number no longer appears to the left of the command. In other words, instead of the following:

  ```
  HP9300(config)#[ 1] en
  ```

  You would see this when you press Ctrl-N, Ctrl-P, or the up and down arrows:

  ```
  HP9300(config)# en
  ```

- If you enter an invalid command followed by ?, a message appears indicating the command was unrecognized. For example:

  ```
  HP9300(config)# rooter ip
  Unrecognized command
  ```

- The Ctrl-Z key now serves a similar function to the **end** command. When you press Ctrl-Z at any CONFIG level of the CLI, activity is moved to the privileged EXEC level. When you press Ctrl-Z at the privileged EXEC level, activity is moved to the user EXEC level.

Software Fixes

This section lists the problems that have been fixed in switch and router software release 06.6.05. For information about fixes in a previous software release, see the release notes for that release.

- **AppleTalk** – The HP routing switch could sometimes merge separate AppleTalk cable ranges sent by some older third-party AppleTalk routers. For example, if such a router sent ranges 100-200 and 3000-3100, the HP routing switch merged them into the single range 100-3100, causing the HP routing switch's AppleTalk route table to be incorrect.
- **AppleTalk** – The filtering option to deny additional zones on an AppleTalk interface also denied the zones configured on that interface, unless they were explicitly permitted. In software release 06.1.09, the option denies all zones except the zones configured on the AppleTalk interface itself.
- **AppleTalk** – The routing switch dropped packets instead of forwarding them if the routing switch did not know the ARP address of the destination. As result, Apple computers sometimes could not see other devices **SRP**

– In a specific active-standby SRP configuration with virtual interfaces, if the network cable from the active router was disconnected and the client cleared its ARP entry for the virtual interface, IP pings to the virtual router IP address failed.

- **AppleTalk** (6308M-SX and 6208M-SX only) – The filtering option to deny additional zones on an AppleTalk interface also denied the zones configured on that interface, unless they were explicitly permitted. In software release 06.6.05 and later, the option denies all zones except the zones configured on the AppleTalk interface itself.
- **AppleTalk** – The router did not broadcast packets that had a MAC broadcast address (09-00-07-ff-ff-ff) and came in through a virtual port.
- **AppleTalk** (9304M and 9308M routing switches with 16-port 10/100 management module only) – Forwarded AppleTalk packets were corrupted.

NOTE: This problem did not affect routing switches using types of management modules other than the 16-port 10/100 module.

- **ARP** (9304M, 9308M, and 6308M-SX) – The ARP address corresponding to a deleted MAC address was not also deleted.
- **ARP** (9304M and 9308M routing switches only) – When the routing switch was forwarding IP packets with a Router Alert option, if the destination IP address was directly connected and the ARP mapping for it was not available initially but was resolved later, the routing switch incurred buffer loss.
- **ACL** – The HP devices did not support named ACLs for SNMP, Web, or Telnet access.
- **Authentication** – If you entered the Telnet password at the Login Name prompt, the software granted access regardless of the authentication method configured for Telnet access.
- **BGP4** – In some cases, a regular expression used to select BGP routes did not work properly when a BGP4 route's AS-path attribute was empty.
- **BGP4** – If a routing switch advertised only a single route in a BGP4 update packet, and if that route was a default route (0.0.0.0), the update packet could sometimes be formatted incorrectly. This could cause corruption in the BGP4 route table on a neighbor that received the corrupted packet.
- **BGP4** – Occasionally, if the routing switch lost a neighbor connection, the CPU would run at a constant utilization rate of 55%.
- **BGP4** – If a static default route was deleted, the router did not automatically learn the default route through BGP4. The default route learned through BGP4 was present in the BGP4 route table but was not imported to the IP route table.
- **BGP4** – If BGP4 learned a default route and the connection to the neighbor router was broken, the routing switch did not learn the default route through OSPF from its OSPF neighbor.
- **BGP4** – In networks containing OpenBSD routers, the routing switch sent a NOTIFICATION error message if it received an UPDATE message from an OpenBSD router and the attribute flag for the partial bit in the message was set to 1.
- **BGP4** – The default maximum number of BGP4 neighbors on Redundant Management modules with 128MB has been changed from 15 to 14.
- **BGP** – The regular expression "_" (underscore) was not supported, even though it was documented. This regular expression is now supported. For usage information for all the regular expressions, see the "Using Regular Expressions" section of the "Configuring BGP4" chapter in *Book 2: Advanced Configuration and Management Guide.* This section also contains corrections to some of the examples for various expressions.
- **BGP** (9304M and 9308M) – In configurations where BGP has multiple paths to a destination but load sharing is not enable the following issues could occur:
  - When evaluating the paths to select the best one, the software might use a null0 static route to derive the next hop for the path.
  - The software could sometimes select a path through a non-direct route instead of a path through a direct

route.

- When updating the path, the software did not update the MAC address for the next hop. As a result, traffic was not forwarded to the appropriate destination.

- **BGP** – The regular expression "_" (underscore) was not supported, even though it was documented. This regular expression is now supported. For usage information for all the regular expressions, see the "Using Regular Expressions" section of the "Configuring BGP4" chapter in *Book 2: Advanced Configuration and Management Guide.* This section also contains corrections to some of the examples for various expressions.
- clear statistics ethernet command – This command did not clear the hardware counters.
- **CLI** – The CLI interface did not properly ignore spaces, such as leading spaces in front of a command. This applied to CLI sessions on the serial port and through Telnet.
- **CLI** – In some situations, a debugging message could appear on the console. Although unexpected, these messages did not indicate system problems.
- **CLI** – If you configure multiple areas, then delete one of the areas, the **show ip ospf config** command did not display any of the areas.
- **CLI** – In some circumstances, certain sequences of keys pressed while in the CLI's CONFIG mode could cause invalid entries in the startup-config file, a system reset, or both.
- **CLI** – The **show ip ospf interface** <*IP-addr*> command displayed Designated Router (DR), Backup Designated Router (BDR), and Neighbor counts incorrectly even though the **show ip ospf interface** command displayed the counts correctly.
- **CLI** – The **show ip ospf neighbor** command lists the wrong port numbers for the links with OSPF neighbors.
- **CLI** (9304M and 9308M only) – When the **no chassis trap-log ps1** command was entered to disable the SNMP trap for failures of power supply 1, then the Syslog buffer was displayed, the buffer contained a message that the power supply had failed, although the power supply had not failed.
- **CLI** - In a situation where a routing switch had learned a default route through OSPF, and you added a static default route to the router, the **show ip route** command display looped at the end of the display.
- **CLI** – A cosmetic error incorrectly displayed the port speed in the CLI prompt for a Gigabit port. This was observed on a 9304M or 9308M routing switch for Gigabit port 8. The CLI prompt was "`HP9300(config-if-e100-1/8)#`" but should have been "`HP9300(config-if-e1000-1/8)#`.
- **CLI** – If the device is configured to authenticate user access, after a user enters a password to start a CLI session, the system name in the command prompt appears twice in each prompt (for example, `HP9300HP9300>`). This is a cosmetic issue only and does not affect the device's operation or performance.
- **CLI** – The **show cpu** command sometimes displays incorrect statistics.
- **CLI** – If the command **no spanning-tree single** was entered when single-instance spanning tree was not enabled, the software removed the VLAN-related configuration from the running-config.
- **CLI banners** – The software placed users who logged in using Telnet into the banner editing mode.

NOTE: The CLI contains a change to the syntax for configuring CLI banners. The **exec_mode** parameter is changed in software release 06.6.05 to **exec**. The **exec_mode** form is supported for backward compatibliity, but when you save configuration information to the startup-config file, the parameter appears as **exec**, even if you enter **exec_mode**. (This change is reflected in the description of the feature in "Configurable CLI Banners" on page 16.

- **DHCP** – (9304M and 9308M, and 6308M-SX routing switches only) DHCP packets could be corrupted while passing through the switch. This problem occurred when the client and server were both in the same VLAN and the VLAN had a virtual interface.
- **DHCP** – If the IP address in a DHCP OFFER packet was previously owned by another device (another MAC address), the router incorrectly forwarded the DHCP reply to that MAC address instead of the device tat sent

the DHCP request.

- **IP** – If a device's IP cache became full, old entries did not age out properly. In some cases, this caused a system reset.
- **IP (routing switches only)** – If the device was booted with a blank startup-config file, the IP directed broadcast feature was enabled on the device. The feature should have been disabled by default.
- **IP ACLs** – IP ACLs configured to permit packets for UDP port 53 (DNS) did not permit the packets.
- **IP Cache** – When a routing switch had a large number of IP cache entries (in excess of 8000), it did not always forward properly to those destinations.
- **IP Multicast** – In a configuration using tagged ports and virtual interfaces, the multicast group was not registered correctly on the virtual interface. As a result, multicast packets were not forwarded in one-armed router configurations.
- **IP Multicast** – This software release fixes the following IP multicast issues:
  - IGMP – Some third-party IP multicast routers did not accept Graft Acknowledgement Received messages sent by HP devices when the third-party routers were checking for upstream neighbors. The HP devices operated within the RFC specifications, but the software has been modified to work with these third-party routers.
  - IGMP snooping – When two IGMP group members joined on an HP device on adjacent ports, sometimes one of the group members was not added to the multicast forwarding database.
  - PIM-DM – Occasionally, the IGMP reports did not remove the prune state, resulting in a delayed join. In this case, the HP device sometimes did not remove the prune state if the reports came in very closely after the data packet.
- **IP** – Some historic reserved IP sub-nets, such as 128.0.x.x/16, were not supported.
- **IP (6308M-SX)** – The device did not properly send or receive IP pings longer than 1473 bytes.
- **IP and trunk groups (6308M-SX)** – If ports were configured in a server trunk group, IP cache entries learned on the ports were not properly cleared from the IP cache when the cache was flushed. This could result in some packets not being forwarded to specific IP addresses.
- **IP (6308M-SX)** – If the device was booted with a blank startup-config file, the IP directed broadcasts feature was enabled on the device. The feature should have been disabled by default.
- **IP (6308M-SX)** – The device did not properly send or receive IP pings longer than 1473 bytes.
- **IP and trunk groups (6308M-SX)** – If ports were configured in a server trunk group, IP cache entries learned on the ports were not properly cleared from the IP cache when the cache was flushed. This could result in some packets not being forwarded to specific IP addresses.
- **IP (6308M-SX)** – If the device was booted with a blank startup-config file, the IP directed broadcasts feature was enabled on the device. The feature should have been disabled by default.
- **IP policies** – The software allowed you to configure a policy to deny all traffic and create SNMP traps and log entries for denied packets, but the software did not allow you to apply the policy to interfaces. For example, the CLI accepted ip policy 1 deny any any log and allowed the command to be saved to the startup-config file, but did not allow you to apply the policy to interfaces. Instead, the software displayed a message such as “Error - filter 1 is not configured in the global table”.
- **IP policies (9304M and 9308M)** – Applying 256 or more IP policy groups to a port or virtual interface was not possible.
- **IPX** – If the routing switch received an IPX unicast packet for a MAC address that the device had not yet learned, the device incorrectly forwarded the packet instead of flooding the packet to the destination VLAN.
- **IPX** – If the link on a port associated with a virtual interface went up or down, the virtual interface stopped accepting IPX packets.
- **IPX** – Forwarding filters did not work properly when filtering on directly connected networks.
- **IPX** – In configurations involving trunk ports and virtual interfaces, IPX broadcast requests from a client were

not forwarded to other VLAN ports.

- **IPX** – Due to back-to-back periodic SAP/RIP updates, slower IPX routers such as NT running Netware Gateway Service sometimes did not learn all the RIP routes and SAP advertisements being broadcast to them.
- **Layer 2 forwarding** (9304M, 9308M, and 6308M-SX only) – In some cases, Layer 2 forwarding performance could be slowed due to unavailability of MAC-related system resources.
- **MAC** (9304M, 9308M, and 6308M-SX only) – A learned MAC address was not relearned if the address moved to another port.
- **Mtrace** – If an Mtrace query is received with a large packet, resulting in one or more fragmented packets, software releases prior to 06.6.05 did not process the fragmented Mtrace packets properly, causing the system to reset.
- **Null0 static routes** – If you removed a null0 static route from the IP route table, the software did not clear the IP cache.
- **OSPF** – If the routing switch had a single neighbor that was a DR-OTHER (other Designated Router) and that neighbor went down, the HP routing switch did not update its LSA accordingly.
- **OSPF** – On rare occasions, deleting an OSPF area caused a checksum failure in the External LSA database and sometimes caused the device to reset.
- **OSPF** – If the routing switch had a single neighbor that was a DR-OTHER (other Designated Router) and that neighbor went down, the HP routing switch did not update its LSA accordingly.
- **OSPF (9304M and 9308M)** – The software set the destination IP address for OSPF routes to match the route's LSA ID, if the route was an External route, before placing the route in the IP route table. In cases where a third-party vendor's OSPF router supported Appendix E in RFC 2338 and thus set the host bits in the LSA ID in External LSAs, the routes contained the host bits in the destination IP address when the HP routing switch placed the routes in the IP route table.

  In software release 06.6.05 and later, the software checks OSPF external routes to determine whether they are host routes. If a route is a host route (with network mask 255.255.255.255), the software sets the host bits before placing the route in the IP route table. Otherwise, the software does not set the host bits in the route before placing it in the IP route table.
- **OSPF** – High-priority OSPF control packets were not forwarded properly within virtual interfaces.
- **OSPF** – IP load sharing on an IP default network did not work.
- **OSPF** – In some instances, when multiple paths existed to the same destination, the software did not use the most specific path.
- **OSPF** – Summary LSAs were generated for networks within the same area.
- **OSPF** – Under some conditions, summary LSAs were generated every 20 seconds.
- **OSPF** – Under some conditions, static routes redistributed through OSPF intermittently got lost when there were changes in the OSPF routes.
- **OSPF** – When an OSPF route lost a path, the IBGP routes that depended on this OSPF route were not updated to use a different path.
- **PIM** – If a MAC address aged out or the MAC table was cleared while a multicast group session was active, the router duplicated one of the multicast packets and forwarded it to the destination. This caused one packet to be sent by the server but two to be received by the client.
- **PIM** – The mechanism used to determine which virtual interfaces have PIM traffic was reporting incorrect active interfaces. The mechanism has been changed so the correct interfaces will be reported active to the software.
- **PIM** – Multicast packets from a server passing through a router were sent out of the router with the source MAC address unchanged. The packets now are changed so that the source MAC address is the router interface’s MAC address.

- **PIM** – In PIM dense mode, in a topology with multiple equal-cost paths to a destination, the forwarding cache was not set correctly. This caused duplication of multicast forwarded packets.
- **PIM** – In PIM dense mode, prune messages were not being processed correctly if the HP device happened to be the designated router for the sub-net.
- **Protection against Denial of Service (DoS) attacks** - This software release fixes the following vulnerabilites to DoS attacks:
  - A DoS attack with a TCP port scan reported all ports on the device were active.
  - A DoS attack that contained an invalid setting for the TCP option key caused HP devices to reset.

NOTE: In addition to these fixes, software release 06.6.05 contains configurable DoS prevention features. See “Protection Against Denial of Service (DoS) Attacks” on page 11.

- **RADIUS support** – (9304M, 9308M, and 6308M-SX routing switches) If an HP device received a RADIUS packet with an extended format, the device reset.
- **Redundant Management modules (9304M and 9308M routing switch only)** – If the active management module was removed from a chassis containing two redundant management modules, the standby module did not regard the removal as a switch-over but instead reloaded the software.
- **RIP** – On routing switches configured to run RIPv1, RIP advertised the same route twice in different packets instead of advertising the route information in one summarized route.
- **RIP** – For routes that were redistributed from OSPF to RIP, if there were two paths to a network, and one of the paths became unavailable, the routing switch sent a RIP packet with a metric of 16, indicating the network was unreachable.
- **Route table (9304M and 9308M routing switch only)** – In some situations, overlapping prefixes in the route table could cause buffer loss on the management module.
- **SNMP** – If the Spanning Tree Protocol (STP) was disabled in the device, the dot1dBaseNumPorts object incorrectly reported that the number of ports was 0. The object now correctly reports the number of ports even if STP is disabled.
- **SNMP** – SNMPv2 GetBulk requests to the SnSwPortDesc object in the port table returned information for only the last port in the device, instead of returning information for all the ports.
- **SNMP** – The compatibility checking for SNMP requests did not allow SNMP packets from third-party SNMP applications if those packets were SNMPv1 packets but contained imbedded 64-bit-counter requests. The SNMP RFCs specify that only SNMPv2 uses 64-bit counters. However, to accomodate the third-party SNMP applications that use 64-bit-counter requests in SNMPv1 packets, the HP software has been changed to allow SNMPv1 packets containing 64-bit counter requests.
- **Static route redistribution** – If a router had a static route with a next hop address on an OSPF network, and the router also had a connection to an OSPF neighbor that was redistributing static routes through OSPF, the router did not redistribute its own static route into RIP after the OSPF interface was brought down and then brought back up.
- **system-max commands** – (9304M, 9308M, and 6308M-SX routing switches) The **system-max mac-filter-sys** and **system-max mac-filter-port** commands did not increase the system capacity for filters. The CLI accepted these commands but the software did not actually change the corresponding table.
- **show tech-support command (9304M and 9308M routing switch only)** – If multiple redundant management modules have been installed, this command could cause the system to reset.
- **Single STP and Fast Span (9304M, 9308M, and 6308M-SX)** – The software sometimes did not "fast" age out the MAC entries in all VLANs. As a result, the MAC and ARP entries were not aged out quickly.
- **SRP** – In a specific active-standby SRP configuration with virtual interfaces, if the network cable from the active router was disconnected and the client cleared its ARP entry for the virtual interface, IP pings to the virtual router IP address failed.
- **SRP** – In some configurations, the source IP address of an SRP virtual router was randomly configured as either the virtual MAC address or the port's physical MAC address, instead of always being the virtual MAC

address. This occurred when another IP interface on the same device was not part of the SRP virtual interface but was actively communicating and was sending gratuitous ARP requests.

- **SRP** – When a learned MAC entry was deleted, the VLAN ID was not checked.
- **SRP** – If you configured a second instance of SRP, the software did not enable the second VIP. As a result, the second SRP instance was not operational.
- **SRP** – Dynamically creating a second instance of SRP using a newly created secondary IP address entry did not take effect until the change was saved to flash memory and the software was reloaded.
- **TACACS+ authentication** – If the HP device was configured to use a TACACS+ server for authentication, but that server was not available, the CLI allowed a user access based on the Enable password. This behavior was accepted based on the authentication-method list for CLI Enable access. The user was then able to proceed to the CONFIG mode. However, the TCP timer used by the CLI for the TACACS+ authentication still ran in the background. When the timer expired, the CLI put the user back in the User EXEC mode, as though the TACACS+ authentication had resulted in denial of access to the Privileged EXEC (Enable) mode. Note that TACACS (as opposed to TACACS+) authentication uses UDP instead of TCP and was not affected by this problem.
- **Traceroute** – The software stopped the trace if an intermediate hop failed to respond. In the current software release, the trace continues even if an intermediate hop fails to respond, until the destination is reached or the specified number of hops has been tried. (The default number of hops is 30.)
- **TACACS** – If a user entered an invalid username or password, the software allowed only one or two seconds to re-attempt the login, then moved to the secondary authentication method. The software now allows more time for re-attempting a login.
- **TACACS/TACACS+ authentication** – If you entered the Telnet password at the Password prompt instead of the password configured on the TACACS/TACACS+ server, the software still granted access to the CLI.
- **TCP sequence numbers** – The Initial Sequence Number (ISN) sent by the HP device in a TCP SYN ACK packet (in response to a TCP SYN packet) was incremented by 1 for each TCP connection to the management IP address, making ISNs generated by the HP device predictable. Such predictability made the HP device vunerable to sequence number attacks, where an attacker could guess the ISN sent by the HP device in response to a TCP SYN packet from a trusted machine. After guessing the ISN, the attacker could then impersonate a connection from the trusted machine to the HP device's management IP address, and potentially execute commands on the HP device.

  In release 06.6.05, TCP sequence numbers are generated randomly for each connection, making it very difficult for an attacker to guess ISNs sent by the HP device.
- **TCP sequence numbers** – If, due to a wraparound in the sequence number counter, the sequence number in an ACK packet was smaller than the last unacknowledged sequence number, the ACK packet was treated as a duplicated ACK, and the last unacknowledged sequence number remained unacknowledged.
- **VLAN** (9304M or 9308M with Redundant Management module only) – You could not actually configure 2048 VLANs on a 9304M or 9308M chassis containing Redundant Management modules.
- **VLAN** (6308M-SX and 6208M-SX) – The system reset if you added two Gigabit ports as tagged ports to a VLAN.
- **VRRP** – When a learned MAC entry was deleted, the VLAN ID was not checked.
- **VRRP** – In some configurations, the backup router could have erroneous MAC CAM entries, which prevented proper failover.
- **VRRP** – In some configurations, using the same VRID on multiple ports for different IP interfaces caused duplicate virtual MAC CAM entries. Software release 06.1.11 provides error checking to ensure against such configurations.

NOTE: VRRP allows up to 255 VRIDs (1 – 255). HP recommends that you do not use the same VRID for different interfaces.

- **VRRP** – In circumstances in which track ports for both routers in a VRRP configuration went down, the Master

Router status was not returned to the VRID owner.

- **VRRP** (chassis routing switches only) – If a backup VRRP router defined on a virtual interface transitioned from Master to Backup, while the "owner" of the interface was becoming the Master again due to track port recovery or a non-abdicate command entered by the user, the Backup sent a gratuitous ARP request to the shared hub or switch. As a result, VRRP clients would continue to send unicast packets to the Backup rather than the Master.
- **VRRP** – If you entered the **backup priority** <*num*> **track-priority** <*num*> command, the software did not write the **priority** command to the startup-config file.
- **VRRP** – Backup VRRP routers received keepalive messages that were intended only for the master VRRP router. As a result, the backup VRRP routers became master routers.
- **Virtual Interfaces** – (6308M-SX) If you sent a ping from one virtual interface to another, then cleared the MAC table, pings did not work properly. In addition, the device sometimes did not relearn source MAC addresses and dropped traffic from those addresses.
- **Web management interface** – The Web management interface did not place TACACS+ authentication parameters into effect until the startup-config file was saved and the device was rebooted, even though the interface accepted the parameters.
- **Web management interface** (9304M, 9308M, and 6308M-SX) – The interface numbers displayed by the Monitor->OSPF->ABR ASBR Routers option of the Web management interface were one number higher than the actual interface numbers. For example, virtual interface “v5” would be listed as “v6”.
- **Web management interface** – Virtual interfaces could not be configured using the Web management interface.
- **Web management interface** – If the device has a virtual interface, the interface tables in the Web management interface listed all possible virtual interfaces, even if not all of those interfaces were configured.
- **Web management interface** – The interface did not properly terminate connections with Lynx (text) browsers.

Known Issues

- **CLI** – The **show ip ospf neighbor** command lists the wrong port numbers for the links with OSPF neighbors.
- **CLI** – If the device is configured to authenticate user access, after a user enters a password to start a CLI session, the system name in the command prompt appears twice in each prompt (for example, `HP9300HP9300>`). This is a cosmetic issue only and does not affect the device’s operation or performance.
- **CLI** (9304M, 9308M, and 6308M-SX) – The **show cpu** command sometimes displays incorrect statistics.
- **CLI** (6308M-SX and 6208M-SX) – If a Gigabit port is configured for Auto-Gigabit (auto-negotiation) in software release 06.6.05 or later, the device does not retain the change following a software reload. The running-config file contains a Negotiation-Off setting for the port and the port at the other end of the link does not come up.
- **DVMRP** – In some configurations, if you run DVMRP with multiple tagged trunk groups, the multicast packets are dropped when a new trunk member port is added to the trunk group.
- **Syslog** – If the link state changes for a port in a trunk group and the port is not the primary port for the group, the software does not send a link state change message to the Syslog buffer.

The information contained in this document is subject to change without notice.

HP Part Number: 5969-2346
Edition 2, March 2000
Printed in USA

HEWLETT® PACKARD