Release Notes *for* Version 07.1.24 Operating Systems

for the HP ProCurve Routing Switch 9304M and 9308M with Redundant Management (MII and MIV)

Software release 07.1.24 supercedes earlier software releases in the 07.x software branch. (For more on software branches, see “Software Branches” on page 5. The 07.1.xx releases are used as follows:

S/W Version:	HP ProCurve 9304M and 9308M Routing Switch Modules:
07.1.10	These Redundant Management Modules: • J4845A ProCurve 9300 GigLX Redundant Management Module (8-port, MII) • J4846A ProCurve 9300 GigSX Redundant Management Module (8-port, MII) • J4847A ProCurve 9300 Redundant Management Module (0-port, MII)
07.1.19 and Higher	These Redundant Management Modules: • All of the modules listed for release 7.1.10. • J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module (8-port, MIV)

These release notes:

- *Summarize* the new operating system enhancements available in software releases 07.1.10 through 07.1.24.
- *Describe* the new operating system enhancements available in software releases 07.1.24.
- *Summarize* earlier software operating problems fixed in software releases 07.1.10 through 07.1.24.

Descriptions of the enhancements in release 07.1.10 are included in the manuals for the 06.6.*xx* and 07.1.*xx* releases. If you purchased a Redundant Management module with software version 07.1.10 or greater installed, then the CD shipped with the module includes these manuals. Otherwise, you can download PDF versions of the latest manuals by visiting **http://www.hp.com/go/hpprocurve** and going to the **technical support** | **manuals** area.

NOTES:

Mini-GBIC ports: Hewlett-Packard offers and supports only mini-GBICs that include an HP label (with product number J4858A or J4859A) for use with the J4856A HP Procurve 9300 Mini-GBIC Module and the J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module. Use of other brands of mini-GBICs is not supported.

Flash Images: The flash image files for this software release differ depending on the product. See “Software Image Files” on page 11.

SNMP: Beginning with software release 05.2.16, 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.

SSH: Beginning with software release 7.1.10, HP supports Secure Shell (SSH) version 1.

Devices Without Redundant Management: For information on upgrading the software on the 9304M and 9308M routing switches WITHOUT redundant management, the 6308M-SX routing switch, and the 6208M-SX switch, see the latest 6.6.x release notes. (Visit the **technical support** | **manuals** area of the HP ProCurve website at **http://www.hp.com/go/hpprocurve**.)

Publication Number

5969-2390
Edition 1
September 2001

Applicable Product

HP Procurve 9304M Routing Switch (J4139A)
HP Procurve 9308M Routing Switch (J4138A)

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are registered trademarks of Microsoft Corporation. Internet Explorer is a trademark of Microsoft Corporation. Ethernet is a registered trademark of Xerox Corporation. Netscape is a registered trademark of Netscape Corporation.

Disclaimer

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

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Hewlett-Packard Company
8000 Foothills Blvd.
Roseville, CA 95747-5551
USA
http://www.hp.com/go/hpprocurve

Contents

— *CONTINUED* —

Software Branches

Beginning with the software releases 06.6.28 and 07.1.10, HP offers two software (Operating System) branches:

- **06.6.28 and later 06.X releases:** These releases typically include only bug fixes, and operate on the following devices:
  - HP 9304M and 9308M routing switches *without* redundant management (that is, with MI modules)
  - HP 6308M-SX routing switch
  - HP 6208M-SX switch
- **07.1.10 and higher 07.1.*X* releases:** These releases typically may include new features, enhancements to existing features, and bug fixes. They operate only with the HP 9304M and 9308M routing switches WITH redundant management (MII or higher modules).

Note Regarding Removing Chassis Modules

When you remove a module from a 9304M or 9308M chassis, disable the module first before removing it from the chassis. Disabling the module before removing it prevents a brief service interruption on other forwarding modules. The brief interruption can be caused by the chassis reinitializing other modules in the chassis when you remove an enabled module.

To disable a module, enter a command such as the following at the Privileged EXEC level of the CLI:

```
HP9308# disable module 3
```

This command disables the module in slot 3.

Syntax: disable module <slot-num>

The <slot-num> parameter specifies the slot number.

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.

NOTE: If you remove the module without first disabling it, the chassis re-initializes the other modules in the chassis, causing a brief interruption in service after which the chassis resumes normal operation.

If, after disabling a module, you decide not to remove the module, re-enable the module using the following command:

Syntax: enable module <slot-num>

For example, to re-enable a module in slot 3:

```
HP9308# enable module 3
```

NOTE: You do not need to enable a module after inserting it in the chassis. The module is automatically enabled when you insert the module into a live chassis or when you power on the chassis.

NOTE: On 9304M and 9308M devices, if you plan to replace the removed module with a different type of module, you must configure the slot for the module. To configure a slot for a module, use the **module** command at the global CONFIG level of the CLI.

Note to IP Multicast Users

HP routing switches support the following IP multicast versions:

- IGMP V2
- PIM Dense mode (PIM-DM) V1
- PIM Sparse mode (PIM-SM) V2
- DVMRP V2

For configuration information, see the “Configuring IP Multicast Protocols” chapter in the *Book 1: Installation and Getting Started Guide*.

Clarification On Trunk Load Sharing

HP devices load share traffic across the ports in a trunk group. The method used for the load sharing depends on the following:

- Device type – 9304M/9308M (chassis) or 6308M-SX and 6208M-SX (fixed-port)
- Traffic type – Layer 2 or Layer 3
- Trunk type – Switch or server
- For certain traffic, port type on which the traffic enters the HP device (Gigabit or 10/100)

NOTE: The port type applies only to Layer 2 traffic on a server trunk group configured on a 9304M or 9308M.

Table 1 lists how HP devices load share traffic across the ports in a trunk group on a 9304M or 9308M.

Table 1: HP Trunk Group Load Sharing – 9304M/9308M

Traffic Layer	Trunk Group Type	Traffic Type	Load-Sharing Basis
Layer 2	Switch	All traffic types	Destination MAC address
	Server	IP received on 10/100 port	Hash value derived from source and destination IP addresses
		IPX received on 10/100 port	Hash value derived from source and destination IPX addresses
		AppleTalk received on 10/100 port	Hash value derived from source and destination AppleTalk addresses
		Other traffic types received on 10/100 port	Hash value derived from source and destination MAC address
		All traffic types received on Gigabit port	Gigabit Port number on which traffic was received
Layer 3	Switch	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address
	Server	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address

Table 2 lists how HP devices load share traffic across the ports in a trunk group on a 6308M-SX or 6208M-SX.

Table 2: HP Trunk Group Load Sharing – 6308M-SX or 6208M-SX

Traffic Layer	Trunk Group Type	Traffic Type	Load-Sharing Basis
Layer 2	Switch	All traffic types	Destination MAC address
	Server	IP	Hash value derived from source and destination IP addresses
		IPX	Hash value derived from source and destination IPX addresses
		AppleTalk	Hash value derived from source and destination AppleTalk addresses
		Other traffic types	Hash value derived from source and destination MAC address
Layer 3	Switch	IP	Source and destination IP addresses
		IPX	Source and destination IPX addresses
		AppleTalk	Source and destination AppleTalk addresses
		Other traffic types	Source and destination MAC address
	Server	IP	Source and destination IP addresses
		IPX	Source and destination IPX addresses
		AppleTalk	Source and destination AppleTalk addresses
		All other	Source and destination MAC address

Redundant Management on the 9304M and 9308M Routing Switches

Redundant Management means that the switch can operate with two management modules installed; one active and one standby. If the active management module becomes unavailable, the standby management module automatically takes over system operation.

Management modules WITHOUT Redundant Management are sometimes termed "MI" modules (for "Management I"). MI modules include:

- J4141A HP ProCurve 9300 10/100 Management Module (16-port)
- J4144A HP ProCurve 9300 Gigabit SX Management Module (8-port)
- J4146A HP ProCurve 9300 Gigabit 4LX/4SX Management Module (8-port)
- J4840A HP ProCurve 6308M-SX Routing Switch

If you are using a management module without redundant management, only one management module can be installed in the routing switch.

Management modules WITH Redundant Management capabilities are sometimes termed "MII" or “MIV” modules (for “Management 2" or “Management 4”). These modules include:

- J4845A ProCurve 9300 GigLX Redundant Management Module (8-port, MII)
- J4846A ProCurve 9300 GigSX Redundant Management Module (8-port, MII)
- J4847A ProCurve 9300 Redundant Management Module (0-port, MII)
- J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module (8-port, MIV)

If you are using a Redundant Management module, you can install either one or two such modules in the routing switch.

NOTE: MI management modules and MII/MIV redundant management modules are mutually exclusive. That is, the routing switch does not operate if an MII or MIV redundant management module is installed while an MI management module is also installed.

For more information, see the chapter titled "Using Redundant Management Modules" in *Book 1: Installation and Getting Started*, that is shipped with your routing switch, available on the CD-ROM included with the routing switch or a management module, and also downloadable from the **technical support** area at **http://www.hp.com/go/hpprocurve**.

These notes also contain information regarding what happens when you disable BGP4, OSPF, or VRRP. See “Usage Guidelines for Access Control Lists (ACLs)” on page 14.

Downloading Software and Documentation

You can download software version 07.1.24 and the corresponding product documentation from HP’s Procurve website as described below.

To Download a Software Version:

1. Go to HP’s ProCurve website at **http://www.hp.com/go/hpprocurve**
2. Click on **software** (in the sidebar).
3. Under “latest software”, click on **switches**.

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:

1. Go to HP’s ProCurve website at **http://www.hp.com/go/hpprocurve**
2. Click on **technical support**, then **manuals**.

3. Click on the name of the product for which you want manuals.
4. On the page listing the manuals, find the latest manuals under the heading "**For software version 06.6.28 and 7.1.10 or greater**".

You will need the Adobe® Acrobat® Reader to view and/or print the manuals.

Software/Device Compatibility

Table 1. Device Compatibility with Software Versions

Device	Supported Software Versions:				
	04791	05084	H2R05216.BIN H2R06605.BIN H2R06616.BIN H2R07110.BIN H2R07122.BIN H2R07124.BIN	HPR05216.BIN HPR06605.BIN HPR06616.BIN HPR06628.BIN HPR06633.BIN HPR06636.BIN	HPS05216.BIN HPS06605.BIN HPS06616.BIN HPS06628.BIN HPS06633.BIN HPS06636.BIN
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) *WITH Redundant Management Module(s) (MII or MIV)*	*No*	*No*	**Yes**	*No*	*No*
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) *WITHOUT Redundant Management Modules (MI)*	**Yes**	**Yes**	*No*	**Yes**	*No*
HP ProCurve Routing Switch 6308M-SX (J4840A)	*n/a*	*n/a*	*No*	**Yes**	*No*
HP ProCurve Switch 6208M-SX (J4841A)	*n/a*	*n/a*	*No*	*No*	**Yes**

If you have a 9304M or 9308M routing switch that was shipped before the software versions described in this document were available, you may want to download either of these releases from HP's ProCurve website. To do so, see the chapter titled "Using Redundant Management Modules" in *Book 1: Installation and Getting Started Guide*, that was shipped with your routing switch or switch.

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 or 9308M with Redundant Management? Here's New Information!

If you received a 9304M or 9308M before software release 07.1.10 began shipping, and you are updating the device to release 07.1.24, then you may want to examine the new product manuals that became available beginning with the 07.1.10 release. To view (and freely download) PDF versions of these manuals (whole manual, or chapter-by-chapter files). See "To Download Product Documentation:" on page 9. (Software features that became available in releases greater than 07.1.10 are described in these release notes.)

Software Image Files

To run either software release 06.6.36 or 07.1.24, you need the indicated boot and flash images listed in the following table.

NOTE: Release 07.1.24 includes Secure Shell (SSH) version 1 (HP 9304M and HP 9308M only).

SSH is not available for the 9304M or 9308M with a 32MB management module (“Management I” module).

Product	Boot Image	Flash Image
HP 9304M HP 9308M With one of these MI modules; that is, WITHOUT Redundant Management: • J4140A • J4144A • J4146A	M1B07108.bin or higher recommended	HPR06636.bin*
HP 9304M HP 9308M With any one or two of these MII or MIV modules; that is, WITH Redundant Management: • J4846A • J4845A • J4847A • J4857A	M2B07108.bin or higher recommended	H2R07124.bin
HP 6308M-SX	• M1B07108.bin or later recommended	• HPR06636.bin*
HP 6208M-SX	• M1B07108.bin or later recommended	• HPS06636.bin*

*These software images do not support Secure Shell (SSH) version 1.

NOTE: If you are adding a Gigabit Copper module to a routing switch chassis (9304M or 9308M), you must upgrade to boot code version M2B07108.bin or later.

Upgrading Software

For easy software image management, all HP devices support the download and upload of software images between the flash modules on the devices and a Trivial File Transfer Protocol (TFTP) server on the network.

The management module contains two flash memory modules:

Primary flash — The default local storage device for image files and configuration files.

Secondary flash — A second flash storage device. You can use the secondary flash to store redundant images for additional booting reliability or to preserve one software image while testing another one.

Only one flash device is active at a time. By default, the primary image will become active upon reload.

You can update the software contained on a flash module using TFTP to copy the update image from a TFTP server onto the flash module. In addition, you can copy software images and configuration files from a flash module to a TFTP server.

NOTE: The 9304M and 9308M routing switches are TFTP clients but not TFTP servers. You must perform the TFTP transaction from the routing switch. You cannot "put" a file onto the routing switch using the interface of your TFTP server.

The 9304M and 9308M TFTP client supports 8 x 3 file names. If you try to copy a file with more than eight characters andu up to three characters in the extension, the interface reports that the file was not found on the TFTP server.

If you are upgrading redundant management modules, the flash code is automatically copied from the active management module to the standby module when you reload. However, the boot code is not automatically copied. See the "Using Redundant Management Modules" chapter in *Book 1: Installation and Configuration Guide* for the routing switches. (For the latest version of this guide, visit **http://www.hp.com/go/hpprocurve** and click on **technical support** | **manuals**.)

Upgrading the Flash Code on Management Modules

When you upgrade the flash code, you must upgrade the flash code on the management module to the same software release ***before*** you reboot.

To upgrade the flash code on a management module:

1. Place the new flash code on a TFTP server to which the routing switch has access.
2. Enter either of the following commands at the Provileged EXEC level of the CLI (example: `HP9300#`) to copy the flash code from the TFTP server into the flash memory of the management module:
   - `copy tftp flash` <ip-addr> <image-file-name> <primary | secondary>
   - `ncopy tftp` <ip-addr> <image-file-name> flash <primary | secondary>
3. Verify that the flash code has been successfully copied by entering the following command at any level of the CLI:

   `show flash`

   The line that begins "`Compressed Pri Code size`" lists the flash code version in the primary flash, at the end of the line. Similarly, the line that begins "`Compressed Sec Code size`" lists the flash code version in the secondary flash.
4. If the flash code version is correct, go to Step 5. Otherwise, go to Step 1.
5. Reload the software by entering one of the following commands:
   - `reload` (This command boots from the default boot source, which is the primary flash area by default.)
   - `boot system flash` <primary | secondary>

For information on changing the block size for TFTP transfers or using the web management interface to transfer software images, refer to *Book 1: Installation and Getting Started Guide* for the routing switches. For an electronic version of the latest release of this guide, visit **http://www.hp.com/go/hpprocurve** and click on **technical support** | **manuals**.)

Updating Boot Code

Under certain conditions, HP support personnel may request you to update the boot code on a routing switch management module. Because the boot code is essential for the management module to operate, and because no backup copy is stored on the module, extreme caution is necessary when updating this code. Use the following steps to verify TFTP operation and to update the boot code.

1. Use the **show flash** command to verify the current boot code version. The last line in this example shows the verification output for boot code version 06.05.00:

```
HP9308> enable
HP9308# show flash
Active management module:
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304, Unit: 2
Boot Flash Type: AMD 29F040, Size: 8 * 65536 = 524288
Compressed Primary Code size = 2579100, Version 07.1.10T53
Compressed Secondary Code size = 2053824, Version 06.6.16T53
Maximum Code Image Size Supported: 2817536 (0x002afe00)
Boot Image Version 06.05.00
```

2. Verify that your TFTP server interoperates properly with the routing switch. To do so, copy the software image stored in secondary flash to a TFTP server, delete the image from secondary flash, and then copy the image you saved onto the TFTP server back into secondary flash. For example, if the IP address of the TFTP server is 192.168.1.1 and the file name you will use to store the image is H2R06616.bin:

   a. Copy the software image stored in secondary flash to the TFTP server.

```
HP9308# copy flash tftp 192.168.1.1 H2R06616.bin sec
HP9308#Flash to TFTP Done.
```

   b. On the routing switch, delete the software image stored in secondary flash and verify that secondary flash is empty. (If secondary flash is empty, you will see "size = 0" in the "Compressed Secondary Code" line of the **show flash** command output.) For example:

```
HP9308# erase flash secondary
Flash Erase HP9308#------------------------------Erase flash Done.
HP9308# show flash
Active management module:
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304, Unit: 2
Boot Flash Type: AMD 29F040, Size: 8 * 65536 = 524288
Compressed Primary Code size = 2579100, Version 07.1.10T53
Compressed Secondary Code size = 0, Version
Maximum Code Image Size Supported: 2817536 (0x002afe00)
Boot Image Version 06.05.00
```

   c. Copy the software image file you just saved (in step a) from the TFTP server back to secondary flash on the routing switch and verify that the code is stored in secondary flash. For example:

```
HP9308# copy tftp flash 192.168.1.1 H2R06616.bin sec
HP9308# Flash Erase ----------------------------------------
Flash Memory Write (8192 bytes per dot)
................................................................................
.....................................................TFTP to Flash Done.

HP9308# show flash
Active management module:
```

```
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304, Unit: 2
Boot Flash Type: AMD 29F040, Size: 8 * 65536 = 524288
Compressed Primary Code size = 2579100, Version 07.1.10T53
Compressed Secondary Code size = 2053824, Version 06.6.16T53
Maximum Code Image Size Supported: 2817536 (0x002afe00)
Boot Image Version 07.01.01
```

The "size" and "Version" values in the "Compressed Secondary Code" line, above, indicate that the software image file has been successfully reloaded into secondary flash. You have now verified that your communication with the TFTP server is working properly.

3. Download the appropriate boot code from the HP Procurve website to your TFTP server. (Go to **http://www.hp.com/go/hpprocurve** and click on **software.**)
4. Use the (undocumented) boot command shown below to initiate the TFTP download. For example, to download the M2B07105.bin boot code from a TFTP server at 192.168.1.1.

CAUTION: It is extremely important that the TFTP download of the boot code is not interrupted. An interruption in this process can result in a non-bootable system. If for any reason the boot code download is not successful, please do not use the **reload** command in the next step. Instead, contact an HP Customer Care Center immediately. To find the HP Customer Care Center for your area, see the support and warranty booklet shipped with your routing switch product, or see the *HP Procurve Networking Service and Support Guide* available on HP's Procurve website at **http://www.hp.com/go/hpprocurve**. (Click on **technical support** and then **support services**.)

```
HP9308# copy tftp flash 192.168.1.1 M2B07108.bin boot
  HP9308# Writing to flash, please wait ... Done

  HP9308# show flash
  Active management module:
  Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304, Unit: 2
  Boot Flash Type: AMD 29F040, Size: 8 * 65536 = 524288
  Compressed Primary Code size = 2579100, Version 07.1.10T53
  Compressed Secondary Code size = 2053824, Version 06.6.16T53
  Maximum Code Image Size Supported: 2817536 (0x002afe00)
  Boot Image Version 07.01.08
```

Note that in the preceding example, the Boot Image Version number (07.01.08) at the end of step 4, above, is a later (higher) version than the Boot Image Version number (06.05.00) at the end of step 1. This indicates a successful download of a new boot image to the `routing switch`.

5. Execute the **reload** command to ensure that the boot code operates properly:

```
HP9308# reload
```

Usage Guidelines for Access Control Lists (ACLs)

This section provides some guidelines for implementing ACLs to ensure wire-speed ACL performance.

For optimal ACL performance, use the following guidelines:

- Apply ACLs to inbound traffic rather than outbound traffic.
- Use the default filtering behavior as much as possible. For example, if you are concerned with filtering only a few specific addresses, create deny entries for those addresses, then create a single entry to permit all other traffic. For tighter control, create explicit permit entries and use the default deny action for all other addresses.
- Use deny ACLs sparingly. When a deny ACL is applied to an interface, the software sends all packets sent or received on the interface (depending on the traffic direction of the ACL) to the CPU for examination.
- Adjust system resources if needed:

- If IP traffic is going to be high, increase the size of the IP forwarding cache to allow more routes. To do so, use the **system-max ip-cache <num>** command at the global CONFIG level of the CLI.
- If much of the IP traffic you are filtering is UDP traffic, increase the size of the session table to allow more ACL sessions. To do so, use the **system-max session-limit <num>** command at the global CONFIG level of the CLI.

Avoid the following implementations when possible:

- Do not apply ACLs to outbound traffic. The system creates separate inbound ACLs to ensure that an outbound ACL is honored for traffic that normally would be forwarded to other ports.
- Do not enable the strict TCP ACL mode unless you need it for tighter security.
- Avoid ICMP-based ACLs where possible. If you are interested in providing protection against ICMP Denial of Service (DoS) attacks, use HP's DoS protection features. See the "Protecting Against Denial of Service Attacks" appendix in *Book 2: Advanced Configuration and Management Guide*. For information on this guide, see "Downloading Software and Documentation" on page 9.

If the IP traffic in your network is characterized by a high volume of short sessions, this also can affect ACL performance, since this traffic initially must go to the CPU. All ICMP ACLs go to the CPU, as do all TCP SYN, SYN/ACK, FIN, and RST packets and the first UDP packet of a session.

ACL Support on the HP Products

HP ACLs have two basic types of uses:

- Filtering forwarded traffic through the device
- Controlling management access to the device itself

In general, routing switches support both types of ACLs. However, the 6208M-SX switch supports ACLs only for access control.

The following table lists the ACL functions supported on each HP routing switch and Layer 2 Switch supported in this software release.

Product	Packet Forwarding ACLs Supported	Management Access ACLs Supported
9304M	Yes	Yes
9308M	Yes	Yes
6308M-SX	Yes	Yes
6208M-SX	No	Yes

Using ACLs and Network Address Translation (NAT) on the Same Interface

You can use ACLs and NAT on the same interface, so long as you follow these guidelines:

- Do not enable NAT on an interface until you have applied ACLs (as described below) to the interface. If NAT is already enabled, you must disable it, apply the ACLs, then re-enable NAT on the interface.
- Enable the strict TCP mode.
- On the inside NAT interface (the one connected to the private addresses), apply inbound ACLs that permit TCP, UDP, and ICMP traffic to enter the device from the private sub-net.

You can use a standard ACL to permit all traffic (including TCP, UDP, and ICMP traffic) or an extended ACL with separate entries to explicitly permit TCP, UDP, and ICMP traffic.

NOTE: You do not need to apply ACLs to permit TCP, UDP, and ICMP traffic unless you are applying other ACLs to the interface as well. If you do not plan to apply any ACLs to a NAT interface, then you do not need to apply the ACLs to permit TCP, UDP, and ICMP traffic.

Here is an example of how to configure a device to use ACLs and NAT on the same interfaces. In this example, the inside NAT interface is port 1/1 and the outside NAT interface is port 2/2.

The following commands enable the strict TCP mode and configure an ACL to permit all traffic from the 10.10.200.x sub-net. A second ACL denies traffic from a specific host on the Internet.

```
HP9308(config)# ip strict-acl-tcp
HP9308(config)# access-list 1 permit 10.10.200.0 0.0.0.255
HP9308(config)# access-list 2 deny 209.157.2.184
```

The following commands configure global NAT parameters.

```
HP9308(config)# ip nat inside source list 1 pool outadds overload
HP9308(config)# ip nat pool outadds 204.168.2.1 204.168.2.254 netmask 255.255.255.0
```

The following commands configure the inside and outside NAT interfaces. Notice that the ACLs are applied to the inbound direction on the inside NAT interface, and are applied *before* NAT is enabled. In this example, ACL 1 permits all traffic to come into the inside interface from the private sub-net. ACL 2 denies traffic from a specific host from going out the interface to the private sub-net.

```
HP9308(config)# interface ethernet 1/1
HP9308(config-if-e1000-1/1)# ip address 10.10.200.1 255.255.255.0
HP9308(config-if-e1000-1/1)# ip access-group 1 in
HP9308(config-if-e1000-1/1)# ip access-group 2 out
HP9308(config-if-e1000-1/1)# ip nat inside
HP9308(config-if-e1000-1/1)# interface ethernet 2/2
HP9308(config-if-e1000-2/2)# ip address 204.168.2.78 255.255.255.0
HP9308(config-if-e1000-2/2)# ip nat outside
```

Where to Find More Information

The following topics are found in *Book 1: Installation and Reference Guide* (for software releases 6.6.x and 7.1.x). For more information, see “Downloading Software and Documentation” on page 9.

- For traffic filtering ACLs, see the “Using Access Control Lists (ACLs)“ chapter in *Book 1: Installation and Getting Started Guide*.
- For management access ACLs, see the “Securing Access to Management Functions“ chapter in *Book 1: Installation and Getting Started Guide*.
- For DoS protection features, see the “Protecting Against Denial of Service Attacks“ appendix in *Book 1: Installation and Getting Started Guide*.
- For information about IP access policies, see the “IP Access Policies“ section in the “Policies and Filters” appendix of *Book 1: Installation and Getting Started Guide*.
- For NAT configuration information, see the “Network Address Translation“ chapter in *Book 1: Installation and Getting Started Guide*.

Maximum File Sizes for Startup-Config and Running-Config Files

Each HP device has a maximum allowable size for the running-config and the startup-config file. If you use TFTP to load additional information into a device's running-config or startup-config file, it is possible to exceed the maximum allowable size. If this occurs, you will not be able to save the configuration changes.

The following table lists the maximum size for the running-config and the startup-config file on HP devices.

Product type	Maximum running-config and startup-config file sizes[a]
A 9304M or 9308M using Management II or higher	256K
A 9304M or 9308M using Management I	128K
A 6308M-SX or 6208M-SX	64K

a.The running-config and startup-config file can each be the size listed. The maximum size is not the maximum combined size for the running-config and startup-config files.

To determine the size of an HP device's running-config or startup-config file, copy it to a TFTP server, then use the directory services on the server to list the size of the copied file. To copy the running-config or startup-config file to a TFTP server, use one of the following commands.

- Command to copy the running-config to a TFTP server:

  copy running-config tftp <ip-addr> <filename>

- Command to copy the startup-config file to a TFTP server:

  copy startup-config tftp <ip-addr> <filename>

Note Regarding Disabling BGP4, OSPF, or VRRP

You can easily disable a routing protocol using the CLI or Web management interface. However, be careful when you disable BGP4, OSPF, or VRRP. When you disable these protocols, the routing switch removes all the configuration information for the disabled protocol from the running-config. Moreover, when you save the configuration to the startup-config file after disabling one of these protocols, all the configuration information for the disabled protocol is removed from the startup-config file.

The CLI displays a warning message such as the following:

```
HP9300(config-bgp-router)# no router bgp
router bgp mode now disabled. All bgp config data will be lost when writing to
flash!
```

The Web management interface does not display a warning message.

If you have disabled the protocol but have not yet saved the configuration to the startup-config file, you can restore the configuration information by re-entering the command to enable the protocol (ex: **router bgp**), or by selecting the web management option to enable the protocol. If you have already saved the configuration to the startup-config file, the information is gone.

If you are testing a BGP4, OSPF, or VRRP configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup-config file containing the protocol's configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

Summaries of Enhancements

This section summarizes the operating system enhancements in software release 07.1.10 through 7.1.24. These enhancements are not described elsewhere in the product documentation.

NOTES: Releases 07.1.19 and higher support the J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module and the J4856A HP Procurve 9300 Mini-GBIC Module.

If your routing switch management module was shipped with release 07.1.19 or later, then the CD-ROM included with the shipment includes this documentation. Otherwise, refer to “Already Using a 9304M or 9308M with Redundant Management? Here’s New Information!” on page 10 for instructions on downloading the documentation from the web. For information about the fixes in this release, see “Software Fixes” on page 50.

Summary of Enhancements in 07.1.24

07.1.24 Layer 3 Enhancements

Enhancement	Description	See Page
New command to list routes redistributed into OSPF	The **show ip ospf redistribute route** command lists the routes that have been redistributed into OSPF and the source from which the routes were redistributed.	29
New display option for **show ip bgp routes** command	The new **nexthop** <ip-addr> option displays the routes that use the specified IP address as their next hop.	29
New display option for **show ip route** command	The new **none-bgp** option displays the most specific route (the route with the longest matching prefix) that was not learned from BGP4.	29
Changes to **show ip bgp routes detail** display	The display contains new fields and does not display attributes that do not apply to the specified route.	30
New command to list routes redistributed into OSPF	The **show ip ospf redistribute route** command lists the routes that have been redistributed into OSPF and the source from which the routes were redistributed.	29

07.1.24 System-Level Enhancements

Enhancement	Description	See Page
AAA security for commands pasted into a running-config file	If AAA security is enabled on the device, commands pasted into the running-config file are subject to the same AAA operations as if they were entered manually.	31
Configurable holddown timer for SNMP traps	You can change the number of seconds a HP device waits after startup before beginning to send SNMP traps to receivers.	31
Displaying information about individual TCP connections	The **show ip tcp status** command displays detailed information about a specified TCP connection, including the sequence and ACK numbers, window sizes, and available buffer sizes.	27
Displaying internal TCP buffer memory allocation	The **debug ip tcp memory** command causes messages to be displayed when memory is allocated or deallocated to the internal TCP buffers.	27

Enhancement	Description	See Page
Displaying TCP connections	The **show ip tcp connections** command displays information about each TCP connection on the device, including the local IP address, local port number, remote IP address, remote port number and the state of the connection. In addition, the command displays the percentage of free memory for each of the internal TCP buffers.	27
Displaying TCP memory usage	The **show memory tcp** command displays the amount of used and free memory for each of the four internal TCP buffers.	26
Maximum number of local user accounts increased	You can configure up to 32 local user accounts for access authentication. Previous releases support up to 16 local user accounts.	26
Support for **spanning-tree** command in VLAN groups	You can enable or disable the Spanning Tree Protocol (STP) in all the VLANs in a VLAN group by enabling or disabling the feature in the VLAN group itself.	28
The “q” key escapes from a scrolling display	You can use the “q” key to escape from a scrolling display in a CLI session. Previous releases support only Ctrl+c to do this.	29

Summary of Enhancements in 07.1.22

07.1.22 Layer 3 Enhancements

Enhancement	Description	See Page
Enhancements to how BGP4 compares route paths	The default method for comparing paths based on the Multi-Exit Discriminator is changed. The software also provides a new option for modifying the MED comparison method.	31
New BGP4 command for refreshing routes redistributed into BGP4	You can place a redistribution parameter change into effect without clearing the entire BGP4 route table, by instead clearing only the redistributed routes.	33
Change to the default value for the IGMP maximum response time	The default IGMP maximum response time is 5 seconds in software release 07.1.22. In previous software releases, the default is 10 seconds.	33

07.1.22 System-Level Enhancements

Enhancement	Description	See Page
New packet error statistic	The **show statistics** display for individual ports contains a new field that distinguishes normal collisions from late collisions when detected on 10 Mbps ports.	33
Real-time display of Syslog messages on the management console	You can enable real-time display of Syslog messages on a management session. The CLI displays the messages when they are generated, so that you do not need to display the Syslog buffer to see the messages.	34

Enhancement	Description	See Page
New HP MIB objects for ACLs	The HP MIB contains new objects for configuring ACLs and applying them to interfaces.	35
New HP MIB objects for CPU utilization statistics	You can use SNMP to get average CPU utilization statistics for the latest one-second, five-second, and one-minute intervals.	39
Higher maximum number of VLANs allowed in a spanning tree	You can configure up to the maximum number of port-based VLANs allowed on a device to be members of a single instance of the Spanning Tree Protocol (STP). Previously, you could configure up to 128 VLANs to use the same spanning tree.	39
Ability to specify different servers for individual AAA functions	In a RADIUS or TACACS+ configuration, you can designate a server to handle a specific AAA task.	39
Encryption of RADIUS and TACACS keys	When you display the configuration of the HP device, the RADIUS and TACACS keys are encrypted.	40
Ability to enter Privileged EXEC Mode After a Telnet or SSH Login	By default, you enter User EXEC mode after a successful login through Telnet or SSH. Optionally, you can configure the device so that you enter Privileged EXEC mode after a Telnet or SSH login.	41
Telnet/SSH login prompt obtained from TACACS+ server	When TACACS+ authentication is configured, the HP device now obtains both the login prompt and the password prompt from the TACACS+ server.	41
Users prompted for Enable or Line password when TACACS+ server is unavailable	If a user attempts to login through Telnet or SSH, but none of the configured TACACS+ servers are available, the user can be prompted for the Enable or Line password.	41
Command authorization and accounting for console commands	The device now supports command authorization and command accounting for CLI commands entered at the console.	41
Increase in the maximum number of IP routes	A redundant management module can now have up to 256,000 IP routes. Previously, up to 200,000 routes were supported.	41

In addition to the enhancements listed above, software release 07.1.22 also contains a new command for disabling a forwarding module before removing it from a 9304M or 9308M. Disabling the module before removing it can prevent brief service interruptions on other forwarding modules. See “Note Regarding Removing Chassis Modules” on page 5.

Summary of Enhancements in 07.1.19

New Hardware Supported by 07.1.19 (and greater)

Hardware	Description	Location:
New Gigabit module for 9304M and 9308M Chassis devices	The J4856A HP Procurve 9300 Mini-GBIC Module and the J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module both have eight slots for miniature Gigabit Interface Converters (mini-GBICs). 1000BaseSX and 1000BaseLX mini-GBICs are available and can be installed and used in any combination.	See below.

07.1.19 Layer 3 Enhancements

Enhancement	Description	See Page
Support for optimized forwarding on the default route	You can optimize a routing switch's Content Addressable Memory (CAM) to cache multiple forwarding entries for the default route, instead of caching entries for individual destination networks.	42
BGP4 next-hop recursion	You can enable BGP4 to recursively look up a route's next hop until the software finds an Interior Gateway Protocol (IGP) path to the route destination. By default, the software performs only one route lookup, regardless of whether the path to the next hop is an IGP path or a BGP path. Only BGP routes with IGP paths to the next-hop gateway are eligible to be installed in the IP route table.	44
New SNMP objects for PIM Sparse Mode (SM)	You can manage PIM SM through an SNMP application by accessing new PIM SM MIB objects.	46

System-Level Enhancements in 07.1.19

Enhancement	Description	See Page
Authentication encryption for BGP4	By default, the software now encrypts the MD5 authentication strings associated with BGP4 neighbors and neighbor peer groups.	48
TACACS+ Exec authorization supports non-HP A-V pairs	To set a user's privilege level using a TACACS+ server, the HP device can accept either a HP-specific A-V pair or a non-HP-specific A-V pair.	47

Summary of Enhancements in 07.1.10

NOTE: HP does not support the use of release 07.1.10 with the J4856A and J4857A Mini-GBIC modules.

Documentation for Release 07.1.10 Enhancements

The details of these enhancements are included in the manuals provided for release 06.6.xx and 07.1.xx . If your routing switch management module was shipped with release 07.1.19 installed, then the CD-ROM included with the shipment includes this documentation. Otherwise, refer to "Already Using a 9304M or 9308M with Redundant Management? Here's New Information!" on page 10 for instructions on downloading the documentation from the web. For information about the fixes in this release, see "Fixed in 07.1.10" on page 59.

07.1.10 Layer 3 Enhancements

Software release 07.1.10 contains the following Layer 3 enhancements:

Enhancement	Description
Support for up to 10,000 static ARP entries	You can configure a routing switch to support up to 10,000 static ARP entries.

Enhancement	Description
Aggregate default network routes	You can configure a routing switch to aggregate default network routes. This option is useful in environments such as ISPs where the routing switch uses default routes for large numbers of destination hosts.
Host-based IP load sharing for specific destination networks	You can configure a 9304M or 9308M routing switch to perform host-based IP load sharing for specific routes while performing network-based IP load sharing for the other routes.
ICMP Router Discovery Protocol (IRDP) enhancements	IRDP is disabled rather than enabled by default. In addition, you can individually configure IRDP parameters.
Option to disable ICMP redirect	You can disable ICMP redirects on a global or individual port basis.
RIP offset lists	You can add to the metrics of specific inbound or outbound routes.
More flexible IP multicast interface numbering	When you configure PIM or DVMRP on a VLAN's virtual interface, you can use a virtual interface with any valid virtual-interface number. You are no longer restricted to using a virtual interface with a number in the range from 1 – 64.
Hardware forwarding for all fragments of IP multicast packets	You can enable a device to forward all fragments of a multicast packet through hardware. In the previous release, the first fragment of a fragmented IP multicast packet received by the device was forwarded in hardware but the remaining fragments went to the CPU for forwarding.
Multicast Source Discovery Protocol (MSDP)	MSDP allows Protocol Independent Multicast (PIM) Sparse routers to exchange routing information for PIM Sparse multicast groups across PIM Sparse domains. Routers running MSDP can discover PIM Sparse sources that are in other PIM Sparse domains.
Dynamic OSPF memory	The software automatically allocates memory when needed. You do not need to manually configure memory and reload the software.
Support for up to 32 OSPF area ranges in each area	This software release allows up to 32 area ranges in an area. Previous software releases allowed up to four ranges in an area.
Support for up to 25,000 External LSAs	This software release allows up to 25,000 External LSAs on a routing switch. Previous releases allow up to 8000 LSAs.
OSPF group Link State Advertisement (LSA) pacing	The routing switch optimizes OSPF performance by pacing transmission of LSAs. The software sends LSAs in groups at regular intervals to conserve bandwidth, instead of sending them according to individual LSA timers.
External LSA reduction	When multiple ASBRs have equivalent routes to advertise to an external routing domain, the ASBR with the highest router ID actually floods the OSPF AS with the AS External LSAs for the other domain. The other ASBRs flush their equivalent LSAs instead of flooding the OSPF AS with functionally equivalent routes to the ones already advertised.
BGP4 re-advertises BGP routes even when OSPF or RIP routes to the same destination have a lower cost	When a RIP, OSPF, or static route to the same destination as a BGP4 route has a lower administrative distance than a learned BGP4 route, the software installs the route with the lower administrative distance into the IP route table and advertises the BGP4 route. In previous releases, the software did not re-advertise a BGP4 route unless the BGP4 route was in the IP route table.

Enhancement	Description
Redistribution changes take place immediately	Changes to BGP4 redistribution parameters for redistributing routes into BGP4 take effect immediately. In previous releases, the changes sometimes take effect only after you reset the routing switch's neighbor sessions, depending on the changes.
Option to redistribute Internal BGP (IBGP) routes into RIP and OSPF	You can override the default BGP4 protocol behavior and redistribute IBGP routes into RIP and OSPF.
Dynamic BGP4 route refresh	You can dynamically refresh BGP4 routes advertised to or received from a neighbor following a filter change without resetting the BGP4 session with the neighbor.
Change to route map processing of ACL or other filtering deny statements	Route maps that use ACLs for input will not match on values that are denied by the input ACLs, IP prefixes, and so on.
Option to clear BGP4 neighbor sessions based on a specific Autonomous System (AS) number.	You can specify an AS number when clearing BGP4 sessions to clear all sessions for neighbors within a specific AS number.
You can specify a route map name when configuring BGP4 network information	You can set or change BGP4 attributes when creating ("sourcing") a local BGP4 route by associating a route map when configuring BGP4 network information.
Enhancements to **set metric** command in route maps	New options allow you to increase or decrease the metric in a route that matches a route map, or remove the route's metric (remove the MED attribute from the route).
Enhancements to **show ip bgp** commands	This release contains the following enhancements to BGP4 show commands: • Network information – You can display BGP4 network information by specifying an IP address within that network. • Route information – You can display BGP4 route information based on the specific criteria. • Neighbor information – You can display routes received from or advertised to a specific neighbor based on specific criteria.
Enhancement to BGP4 Syslog message	The BGP4 Syslog message for a dropped neighbor session is enhanced to list the reason the session was dropped.
Network Address Translation (NAT)	You can configure an HP device to provide address translation from private addresses to public (Internet) addresses. **Note**: This feature is supported on all chassis routing switches with Management II modules.
Virtual Router Redundancy Protocol Extended (VRRPE)	VRRPE is an extended version of the RFC-standard VRRP that provides the benefits of the RFC-based protocol while also overcoming the protocol's architectural limitations.

07.1.10 Layer 2 Enhancements

Enhancement	Description
Updated STP port Path Cost defaults	The default value for the STP port Path Cost parameter has been changed in accordance with the updated STP specification (IEEE P802.1D). The new value depends on the port speed.
Compatibility with Cisco Systems' Per VLAN Spanning Tree (PVST)	You can enable HP devices to interoperate with devices running PVST.

07.1.10 System Level Enhancements

Enhancement	Description
Enhanced software version information	The **show version** and **show flash** commands provide more information about the software on the device.
New strict mode for ACL processing of UDP traffic	You can configure a HP device to send all UDP packets to the CPU for ACL comparison, instead of just the first UDP packet with specific source and destination information.
New MIB tables for Adaptive Rate Limiting	The HP MIB contains two new tables for port and VLAN Adaptive Rate Limiting information.
Fixed Rate Limiting	You can configure a strict rate limit on a port's inbound or outbound traffic. The device forwards traffic that is within the limit but drops all traffic that exceeds the limit for the specified traffic direction.
Adaptive Rate Limiting	You can configure a flexible bandwidth limit that allows for bursts above the limit, and specify separate actions for conforming and excess traffic.
Denial of Service (DoS) protection for TCP SYN and ICMP transit traffic	You can protect TCP SYN and ICMP traffic being routed by the device against DoS attacks. Previous releases allow you to protect against DoS attacks in traffic addressed to the device itself, but not in traffic the device is forwarding to another device.
Authorization and Accounting support for RADIUS and TACACS+	HP devices now support Authorization and Accounting functions for RADIUS and TACACS+, in addition to the Authentication previously supported.
TACACS+ password prompt support	The TACACS+ password prompt displayed on the device is the one specified by the TACACS+ server.
VLAN-based management access control	You can restrict management access to an HP device to ports within a specific port-based VLAN.
RSA authentication for SSH	HP devices support RSA public-private key authentication for SSH, you can place a list of clients' authorized public keys on the device.
SCP support for secure file transfers	HP devices support Secure Copy (SCP) for securely transferring files to and from remote hosts

Enhancement	Description
Automatic load re-distribution following a healed trunk link	As in previous releases, if a link in a trunk group goes down, the software redistributes the load-balanced traffic across the remaining links. In the current software release, the software also rebalances the traffic when a down link comes back up.
Support for up to 4095 VLANs and up to 4095 virtual interfaces (VEs)	You can configure an HP device to allow up to 4095 VLANs and up to 4095 virtual interfaces. **Note**: This enhancement applies only to the 9304M and 9308M routing switches with Management II or higher modules. The number of VLANs and virtual interfaces supported depends on the amount of DRAM memory on the management module.
VLAN and virtual interface groups	You can simplify configurations that contain many VLANs or virtual interfaces by configuring VLAN or virtual interface groups. A group lets you configure the VLAN or virtual interface attributes one time, then apply the attributes to multiple VLANs or virtual interfaces. **Note**: VLAN groups are supported on the 9304M and 9308M with Management II and higher modules. Virtual interface groups are supported only on the 9304M and 9308M routing switches.
Enhanced CLI for managing redundant management modules	The CLI commands for managing redundant management modules now appear in their own CLI level. This release also contains some new redundant management module commands.
Super Aggregated VLANs	You can configure Layer 2 port-based VLANs within other Layer 2 port-based VLANs. This feature is especially useful for providing each user of a Metropolitan Area Network (MAN) with a private broadcast domain within a larger Layer 2 pipe.
Support for simultaneous Telnet configuration by multiple users	You can enable a device to allow multiple users, on different Telnet CLI sessions, to edit configuration information on the device.
New CLI command for displaying dynamic memory utilization	The **show memory** command display the current utilization of dynamic memory for BGP4 and OSPF.
SNMPv2 view	You can use Access Control Lists (ACLs) to control access to SNMP Management Information Base (MIB) objects on an HP device.
Enhancement to **show default values** command	The **show default values** command now displays the current maximum setting for system parameters, in addition to the default setting and the maximum configurable setting.
CLI enhancements to the startup-config and running-config files	This software release includes the following enhancement to the startup-config and the running-config files: • Route maps are listed right above ACLs, near the end of the file, for easier viewing. • The “permit” parameter in ACLs is spelled fully, rather than abbreviated to “perm”. • The **snmp-server trap-source** command is placed with the other **snmp-server** commands.
Page display is configurable for individual CLI management sessions	Serial console and Telnet CLI users can individually enable or disable page-display mode without affecting the page-display mode of other CLI users.

Enhancement	Description
CLI enhancement to display the idle time for open CLI sessions	The **show who** and **show telnet** commands are enhanced to list the idle time for open CLI sessions.
New CLI command for displaying TACACS+ or RADIUS information	The **show aaa** command displays information about the TACACS+ or RADIUS configuration.
Enhancement to the **show web** command	The **show web** command now displays the privilege level of Web management interface users.
New option for setting the timeout for Telnet sessions	You can change the timeout for Telnet sessions to a value from 1 – 10 minutes.
Enhancements to **show interface** command	The command distinguishes between the down state and the disabled state. In previous releases, the command listed “down” for both states.
ACL configuration supported in the Web management interface	You can configure standard and extended ACLs using the Web management interface.
Greeting banners are displayed at the beginning of a Web management session	If you configure a CLI banner greeting, the greeting also is displayed at the beginning of Web management sessions with the device.
Increasing the Syslog buffer size does not clear entries	You can increase the size of the Syslog buffer without losing the entries that are already in the buffer.
The newline character does not appear in Syslog and SNMP trap messages	Syslog and SNMP trap messages no longer contain a newline character.

Description of Enhancements in 07.1.24

Increased Maximum Number of Local User Accounts

Software release 07.1.24 supports up to 32 local user accounts, used for authenticating management access to the HP device. Previous releases support up to 16 local user accounts.

Displaying TCP Memory Usage

The **show memory tcp** command displays the amount of used and free memory for each of the four internal TCP buffers. For example:

```
HP9308# show memory tcp
TCP MEMORY USAGE
    TCB usage: total=65025, free=63750
    TCP QUEUE BUFFER usage: total=28616, free=18032
    TCP SEND BUFFER usage: total=382500, free=121500
    TCP RECEIVE BUFFER usage: total=382500, free=360000
    TCP OUT OF SEQUENCE BUFFER usage: total=19900, free=19900
```

Syntax: show memory tcp

For each internal buffer, the amount of used and free memory is shown in bytes.

Displaying TCP Connections

The **show ip tcp connections** command displays information about each TCP connection on the device, including the local IP address, local port number, remote IP address, remote port number and the state of the connection. In addition, the command displays the percentage of free memory for each of the internal TCP buffers. For example:

```
HP9308# show ip tcp connections
Local IP address : port  <-> Remote IP address : port  TCP state
TCP: 10.10.10.25 : 23  <-> 10.10.10.15      : 2465  ESTABLISHED
TCP: 10.10.10.25 : 80  <-> 10.10.10.30      : 4026     FIN-WAIT-2
TCP: 10.10.10.25 : 22  <-> 10.10.10.50      : 3578     ESTABLISHED
TCP: 10.10.10.25 : 23  <-> 10.10.10.15      : 2468  ESTABLISHED
TCP: 10.10.10.25 : 23  <-> 10.10.10.15      : 2466  ESTABLISHED
Total 5 TCP connections

TCP MEMORY USAGE PERCENTAGE
FREE TCB = 96 percent
FREE TCP QUEUE BUFFER = 62 percent
FREE TCP SEND BUFFER = 25 percent
FREE TCP RECEIVE BUFFER = 100 percent
FREE TCP OUT OF SEQUENCE BUFFER = 100 percent
```

Syntax: show ip tcp connections

Displaying Information About Individual TCP Connections

The **show ip tcp status** command displays detailed information about a specified TCP connection, including the sequence and ACK numbers, window sizes, and available buffer sizes. For example:

```
HP9308# show ip tcp status 10.10.10.25 23 10.10.10.15 2465
TCP: TCB = 0x210de40a
TCP: 10.10.10.25:23 <-> 10.10.10.15:2465: state: ESTABLISHED
  Send: initial sequence number = 1453320
  Send: first unacknowledged sequence number = 1532710
  Send: current send pointer = 1532710
  Send: next sequence number to send = 1532710
  Send: remote received window = 0
  Send: total unacknowledged sequence number = 3773
  Send: total used buffers 43
  Receive: initial incoming sequence number = 17806845
  Receive: expected incoming sequence number = 17846856
  Receive: received window = 16384
  Receive: bytes in receive queue = 0
  Receive: congestion window = 1460
```

Syntax: show ip tcp status <local IP address> <local port> <remote IP address> <remote port>

Displaying Internal TCP Buffer Memory Allocation

The **debug ip tcp memory** command causes messages to be displayed when memory is allocated or deallocated to the internal TCP buffers. For example:

```
HP9308# debug ip tcp memory
```

For example, when a user establishes a Telnet session with the device, an then terminates it, messages such as the following appear at the destination specified for debugging output. You can turn off these messages with the **no debug ip tcp memory** command.

```
TCP TCB ALLOCATED 210de822
TCP SEND BUFFER ALLOCATED 2111ec80
TCP SEND QUEUE BUFFER ALLOCATED 210d88dc
TCP SEND BUFFER ALLOCATED 2113695c
```

```
TCP SEND QUEUE BUFFER ALLOCATED 210d9714
TCP SEND BUFFER ALLOCATED 2111f838
TCP SEND QUEUE BUFFER ALLOCATED 210d894c
TCP SEND BUFFER ALLOCATED 21117174
TCP SEND QUEUE BUFFER ALLOCATED 210d8444
TCP SEND BUFFER ALLOCATED 210f4aac
TCP SEND QUEUE BUFFER ALLOCATED 210d6fb4
TCP SEND BUFFER ALLOCATED 210f5088
TCP SEND QUEUE BUFFER ALLOCATED 210d6fec
TCP SEND BUFFER FREED 2111ec80
TCP QUEUE BUFFER FREED 210d6fec
TCP RECEIVE QUEUE BUFFER ALLOCATED 210d6fec
TCP RECEIVE BUFFER ALLOCATED 21151530
TCP RECEIVE BUFFER FREED 21151530
TCP QUEUE BUFFER FREED 210d6fec
TCP RECEIVE QUEUE BUFFER ALLOCATED 210d6fec
TCP RECEIVE BUFFER ALLOCATED 21151530
TCP RECEIVE BUFFER FREED 21151530
TCP QUEUE BUFFER FREED 210d6fec
TCP TCB FREED 210de822
```

Syntax: [no] debug ip tcp memory

NOTE: Output from this command appears only on the console or syslog. The output is suppressed when sent to a Telnet or SSH session.

Setting STP State for All VLANs in a VLAN Group

Software release 07.1.24 simplifies configuration of STP in the VLANs in a VLAN group by allowing you to enter the spanning-tree command in the VLAN group itself. All the VLANs in the group inherit the STP state you specify in the group itself.

This feature is especially useful for single-instance STP environments. To configure the VLANs in a VLAN group for single STP, globally enable single STP, then enable STP in the VLAN group.

NOTE: If you do not set the STP state in the VLAN group itself, the STP state of each VLAN depends on the setting in that VLAN.

NOTE: The device has a maximum number of spanning trees that can be configured. If you are not using single-instance STP but you enable STP in a VLAN group, each VLAN you add to the group adds a spanning tree to the configuration. If you add VLANs that cause the maximum number of spanning trees to be exceeded, the VLANs will not work properly.

To enable single-instance STP and set the STP state for all VLANs in a VLAN group, enter commands such as the following:

```
HP9308(config)# spanning-tree single
HP9308(config)# vlan-group 1 vlan 2 to 1000
HP9308(config-vlan-group-1)# tagged 1/1 to 1/2
HP9308(config-vlan-group-1)# spanning-tree
```

These commands enable single-instance STP, then configure a VLAN group consisting of VLANs 2 – 1000, on tagged ports 1/1 and 1/2, and with STP enabled.

Syntax: [no] spanning-tree

New Option for Escaping from a Scrolling Display

In software release 07.1.24, you can use either of the following methods to escape from a scrolling display in a CLI session:

- Press Ctrl+c
- Press q

Support for the “q” key is new in this software release. The Ctrl+c method is supported in previous releases as well as the current release.

Displaying the Routes that Have Been Redistributed into OSPF

To display the routes that have been redistributed into OSPF, enter the following command at any level of the CLI:

```
HP9308# show ip ospf redistribute route
  4.3.0.0 255.255.0.0 static
  3.1.0.0 255.255.0.0 static
  10.11.61.0 255.255.255.0 connected
  4.1.0.0 255.255.0.0 static
```

In this example, four routes have been redistributed. Three of the routes were redistributed from static IP routes and one route was redistributed from a directly connected IP route.

Syntax: show ip ospf redistribute route [<ip-addr> <ip-mask>]

The <ip-addr> <ip-mask> parameter specifies a network prefix and network mask. Here is an example:

```
HP9308# show ip ospf redistribute route 3.1.0.0 255.255.0.0
  3.1.0.0 255.255.0.0 static
```

Displaying BGP4 Routes for a Specified Next-Hop Address

The **show ip bgp routes** command has a new option that enables you to display the routes for a given next-hop IP address. Here is the syntax:

Syntax: show ip bgp routes [detail] nexthop <ip-addr>

If you enter the command without the **detail** parameter, the display has the same fields as other summary **show ip bgp routes** commands. If you use the **detail** parameter, the display has the same fields as other detailed **show ip bgp routes** commands. These fields include:

- Network information — You can display BGP4 network information byspecifying an IP address within that network.
- Route information — You can display BGP4 route information based on the specific criteria.
- Neighbor information — You can display routes received from or advertised to a specific neighbor based on specific criteria.

This section does not describe the **nexthop** <ip-addr> parameter specifically, but the field descriptions are applicable.

Displaying the Longest Matching Non-BGP4 Route in the IP Route Table

The **show ip route** command has a new option, **none-bgp**. This option displays the longest-matching route (the most specific route) that did not come from BGP4, for a specified destination address.

For example, assume that the IP route table has two routes to destination network 100.100.100.1:

- 100.100.100.0/24 – from BGP4
- 100.100.0.0/16 – from OSPF

The IP route table received one of the routes from BGP4 and received the other route from OSPF. The route received from BGP4 is more specific (has a longer matching prefix).

Without the **none-bgp** option, the **show ip route** command displays the most-specific route for the specified network. Here is an example:

```
HP9308# show ip route 100.100.100.1
Total number of IP routes: 37
      Network Address   NetMask           Gateway       Port     Cost      Type
      100.100.0.0       255.255.255.0     10.0.0.1      1/1      1         B
```

Notice that the route learned from BGP4 is displayed, since this route is more specific than the route learned from OSPF.

To display the most-specific route that was not learned from BGP4, enter a command such as the following:

```
HP9308# show ip route 100.100.100.1 none-bgp
Total number of IP routes: 37
      Network Address   NetMask           Gateway       Port     Cost      Type
      100.100.0.0       255.255.0.0       10.0.0.1      1/1      1         O
```

Syntax: show ip route <ip-addr> none-bgp

Changes to the Detailed BGP4 Route Display

The format of the **show ip bgp routes detail** command's output is changed in 07.1.24, and the output has two new fields. Here is an example of the new display:

```
HP9308# show ip bgp routes detail
Total number of BGP Routes: 101863
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
1       Prefix: 3.0.0.0/8,  Status: BE,  Time: 5h16m59s
        NEXT_HOP: 207.69.223.171,  Learned from Peer: 207.69.223.171 (4355)
        LOCAL_PREF: 120,  MED: 10,  ORIGIN: igp,  Weight: 0
        AS_PATH: 4355 701 80
         COMMUNITIES: 4355:65001 4355:65002 4355:65003 4355:65004
          Adj_RIB_out count: 99, Admin distance 40
```

The changes are shown in bold type.

- Time field – This is a new field that indicates the last time an update occurred.
- Adj_RIB_out count field –This field is called RIB_out in earlier software releases.
- Admin distance field – This new field displays the administrative distance of the route.

In addition, if a BGP4 route attribute is not present in the route, the CLI does not display a field for that attribute. For example, the route shown above does not have the CLUSTER_LIST attribute.

Displaying the Routes that Have Been Redistributed into OSPF

To display the routes that have been redistributed into OSPF, enter the following command at any level of the CLI:

```
HP9300# show ip ospf redistribute route
   4.3.0.0.255.255.0.0 static
   3.1.0.0.255.255.0.0 static
   10.11.61.0.255.255.255.0 connected
   4.1.0.0.255.255.0.0 static
```

In this example, four routes have been redistributed. Three of the routes were redistributed from static IP routes and one route was redistributed from a directly connected IP route.

Syntax: show ip ospf redistribute route [<ip-addr> <ip-mask>]

The <ip-addr> <ip-mask> parameter specifies a network prefix and network mask. Here is an example:

```
HP9300# show ip ospf redistribute route 3.1.0.0.255.255.0.0
   3.1.0.0.255.255.0.0 static
```

Setting the SNMP Trap Holddown Time

When an HP device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the device might not be able to reach the servers, in which case the messages are lost.

By default, a HP device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the device sends the traps, including traps such as “cold start” or “warm start” that occur before the holddown time expires.

Software release 07.1.24 enables you to change the holddown time. You can set the holddown time to a value from one second to ten minutes. The default is still one minute.

USING THE CLI

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI:

```
HP9308(config)# snmp-server enable traps holddown-time 30
```

The command in this example changes the holddown time for SNMP traps to 30 seconds. The device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: [no] snmp-server enable traps holddown-time <secs>

The <secs> parameter specifies the number of seconds and can be from 1 – 600 (ten minutes). The default is 60 seconds.

AAA Security for Commands Pasted Into the Running-Config File

If AAA security is enabled on the device, commands pasted into the running-config file are subject to the same AAA operations as if they were entered manually.

When you paste commands into the running-config file, and AAA command authorization and/or accounting is enabled on the device, AAA operations are performed on the pasted commands. The AAA operations are performed before the commands are actually added to the running-config file. The server performing the AAA operations should be reachable when you paste the commands into the running-config file.

NOTE: Since RADIUS command authorization relies on a list of commands received from the RADIUS server when authentication is performed, it is important that you use RADIUS authentication when you also use RADIUS command authorization.

Description of Enhancements in 07.1.22

BGP4 Route Comparison Using the MED

Software release 07.1.22 contains an enhancement to the algorithm the routing switch uses to compare multiple BGP4 paths to the same route destination. The enhancement changes the way the software compares the Multi Exit-Discriminators (MEDs) of otherwise equivalent paths. The MED is the route’s BGP4 metric (cost).

The routing switch uses the following algorithm to compare BGP4 routes. If the algorithm still has more than one path after Step 6, Step 7 compares the MEDs of the paths.

1. Is the next hop accessible though an Interior Gateway Protocol (IGP) route? If not, ignore the route.
2. Use the path with the largest weight.
3. If the weights are the same, prefer the route with the largest local preference.
4. If the routes have the same local preference, prefer the route that was originated locally (by this BGP4 routing switch).
5. If the local preferences are the same, prefer the route with the shortest AS-path. An AS-SET counts as 1. A confederation path length, if present, is not counted as part of the path length.

6. If the AS-path lengths are the same, prefer the route with the lowest origin type. From low to high, route origin types are valued as follows:
   - IGP is lowest
   - EGP is higher than IGP but lower than INCOMPLETE
   - INCOMPLETE is highest
7. If the routes have the same origin type, prefer the route with the lowest MED. The routing switch compares the MEDs based on one or more of the following conditions.
   - By default, the routing switch compares the MEDs of paths ***only if*** the first AS in the paths is the same. (The routing switch skips over the AS-CONFED-SEQUENCE if present.)
   - You can enable the routing switch to also compare the MEDs if the neighbors that advertised the paths are each in the same AS. To enable this comparison, enter the **deterministic-med** command at the BGP4 configuration level of the CLI. This option is new in software release 07.1.22. This option is disabled by default.
   - In addition, you can enable the routing switch to always compare the MEDs, regardless of the AS information in the paths. To enable this comparison, enter the **always-compare-med** command at the BGP4 configuration level of the CLI. This option is disabled by default.

NOTE: HP recommends that you enable the **deterministic-med** option. Enabling this option ensures that the MEDs are compared across all routes received from the same AS.

NOTE: If the path does not have the MED attribute, HP's BGP4 uses zero as the MED value for the comparison.

8. Prefer routes in the following order:
   - Routes received through EBGP from a BGP4 neighbor outside of the confederation
   - Routes received through EBGP from a BGP4 router within the confederation
   - Routes received through IBGP
9. If all the comparisons above are equal, prefer the route with the lowest IGP metric to the BGP4 next hop. This is the closest internal path inside the AS to reach the destination.
10. If the internal paths also are the same, prefer the route that comes from the BGP4 router with the lowest router ID.

NOTE: HP routing switches support BGP4 load sharing among multiple equal-cost paths. BGP4 load sharing enables the routing switch to balance the traffic across the multiple paths instead of choosing just one path based on router ID.

USING THE CLI

In software releases earlier than 07.1.22, the routing switch performs Step 7 as follows:

- The routing switch compares the MEDs of the paths only if the neighbors from which the routing switch learned the paths are in the same AS.
- You can enable the routing switch to always compare the MEDs by entering the **always-compare-med** command. In this case, the routing switch compares the MEDs of the paths regardless of whether the neighbors that advertised the paths are in the same AS.

To enable the routing switch to compare the MEDs only when the neighbors that advertised them are in the same AS, enter the following command at the BGP4 configuration level of the CLI:

```
HP9308(config)# deterministic-med
```

Syntax: [no] deterministic-med

Refreshing Routes Redistributed into BGP4

In software releases earlier than 07.1.22, a change to a redistribution parameter does not affect routes that had already been redistributed into BGP4. The redistribution parameter change itself takes effect immediately, but the change is not applied to the routes that are already in the BGP4 route table. To retroactively place a redistribution change into effect in the earlier software releases, you need to clear the BGP4 route table. This includes clearing the routes that were not redistributed, such as routes learned from BGP4 neighbors.

Software release 07.1.22 enables you to retroactively place redistribution changes into effect without clearing all the routes from the BGP4 route table. You can clear the redistributed routes from the BGP4 route table without affecting the other routes. After you clear the redistributed routes, the software redistributes routes into the table using the redistribution parameters you changed.

USING THE CLI

To refresh the BGP4 route table following a change to a redistribution parameter, enter the following command at the Privileged EXEC or a configuration level of the CLI:

```
HP9308(config)# clear ip bgp local routes
```

Syntax: clear ip bgp local routes

Change to IGMP Maximum Response Time Default

The IGMP maximum response time parameter 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.

In software releases earlier than 07.1.22, the default is 10 seconds. In software release 07.1.22, the default is changed to 5 seconds.

USING THE CLI

To change the IGMP maximum response time, enter a command such as the following at the global CONFIG level of the CLI:

```
HP9308(config)# ip igmp max-response-time 8
```

Syntax: [no] ip igmp max-response-time <num>

The <num> parameter specifies the number of seconds and can be a value from 1 – 10. The default in software release 07.1.22 and later is 5.

New Packet Error Statistic

Software release 07.1.22 contains a new field in the detailed **show statistics** display for an individual port. The new field is called OutLateCollisions. This field is useful for detecting a mismatch in the operating mode (half- or full-duplex) of ports connected to a 10 Mbps or 100 Mbps network segment.

- For 10 Mbps ports, the OutLateCollisions field is incremented each time the port detects an Ethernet collision that occurs 51.2 microseconds or later, after the data was transmitted onto the network segment the port is connected to.
- For 100 Mbps ports, the OutLateCollisions field is incremented each time the port detects an Ethernet collision that occurs 5.12 microseconds or later (one-tenth the interval for 10 Mpbs ports), after the data was transmitted onto the network segment the port is connected to.

A collision that occurs before the late collision time (51.2 microseconds for 10 Mbps and 5.12 microseconds for 10 Mbps) is considered to be a normal collision and is counted in the OutCollisions field instead.

In addition to the new field, the statistics counters for the InErrors and OutCollisions fields have been modified to more accurately distinguish runt packets that are errors from runt packets caused by normal collisions on half-duplex 10 Mpbs ports.

- In software releases earlier than 07.1.22, runts are counted in the InErrors field, for 10 Mbps and 100 Mbps ports, running in half-duplex or full-duplex mode. The runts also are counted in the OutCollisions field, since runts are a common result of a collision.

- In software release 07.1.22 and later, runts are still counted in the OutCollisions field. However, a runt is not counted in the InErrors field for 10 Mpbs ports running in half-duplex mode.
  - For 10 Mbps ports in half-duplex mode, the software does not increment the InErrors field when a runt packet is detected. Instead, the runt is counted in the OutCollisions field.
  - For 100 Mpbs ports, and for 10 Mpbs ports in full-duplex mode, the software does increment the InErrors field when a runt packet is detected.
  - Statistics for runt packets are not applicable to Gigabit Ethernet ports, so the field is not used for this port type.

NOTE: The InCollisions field is not used for any ports and should always contain the value 0.

Here is an example of the new **show statistics** display for an individual port. The new field is shown in bold type.

```
HP9308# show statistics ethernet 2/1
Port  Link State     Dupl Speed Trunk Tag Priori MAC            Name
2/1   Up   Forward   Half 10M   None  No  level0 00e0.52a9.2b00

Port 2/1 Counters:
                InOctets           45994            OutOctets           31528
                  InPkts             562              OutPkts             48
          InBroadcastPkts            462        OutBroadcastPkts           1
          InMulticastPkts             38        OutMulticastPkts           0
               InDiscards              0             OutDiscards           0
                 InErrors              0               OutErrors           0
             InCollisions              0           OutCollisions           0
                                                 OutLateCollisions           0
                Alignment              0                   FCS           0
                GiantPkts              0              ShortPkts           0
             InBitsPerSec           1160           OutBitsPerSec         832
             InPktsPerSec              1           OutPktsPerSec           0
            InUtilization          0.01%          OutUtilization        0.00%
```

Enabling Real-Time Display of Syslog Messages

By default, to view Syslog messages generated by a HP device, you need to display the Syslog buffer or the log on a Syslog server used by the HP device.

Software release 07.1.22 allows you to enable real-time display of Syslog messages on the management console. When you enable this feature, the software displays a Syslog message on the management console when the message is generated.

When you enable the feature, the software displays Syslog messages on the serial console when they occur. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

USING THE CLI

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI:

```
HP9308(config)# logging console
```

Syntax: [no] logging console

This command enables the real-time display of Syslog messages on the serial console. You can enter this command from the serial console or a Telnet or SSH session.

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session:

```
telnet@HP9308# terminal monitor
Syslog trace was turned ON
```

Syntax: terminal monitor

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@HP9308# terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed:

```
telnet@HP9308# terminal monitor
Syslog trace was turned ON
SYSLOG: <9>HP9308, Power supply 2, power supply on left connector, failed

SYSLOG: <14>HP9308, Interface ethernet6, state down

SYSLOG: <14>HP9308, Interface ethernet2, state up
```

New MIB Objects for ACLs

The HP MIB contains the following new tables:

- snAgAclTable – Contains configuration objects for ACLs
- snAgAclBindToPortTable – Contains configuration objects for applying ACLs to interfaces

To use these objects, you need to use HP MIB file MIB07122.mib or later.

snAgAclTable

This table contains the following ACL configuration objects.

Table 3: snAgAclTable Objects

Object	Description
snAgAclIndex	The ACL MIB table index
snAgAclNumber	The ACL ID, which can be in one of the following ranges: • 1 – 99 for standard ACL • 100 – 199 for extended ACL
snAgAclName	The ACL name, if you are using a name instead of a numerical ID to identify the ACL.
snAgAclAction	The action that the software takes if a packet matches the ACL: 0 – deny 1 – permit
snAgAclProtocol	The IP protocol, which can be a number from 0 – 255. The value “0” matches on any IP protocol. **Note**: This object applies only to extended ACLs.
snAgAclSourceIp	The source IP address.
snAgAclSourceMask	The source IP network mask.

Table 3: snAgAclTable Objects (Continued)

Object	Description
snAgAclSourceOperator	The comparison operator, when the source IP protocol is 6 (TCP) or 17 (UDP). The operator can be one of the following: • 0 – equal to • 1 – not equal to • 2 – less than • 3 – greater than • 4 – range • 7 – The operator is not defined **Note**: This object applies only to extended ACLs.
snAgAclSourceOperand1	The first TCP or UDP port number in the source comparison, if applicable. If the value is 0, then this object does not apply to this ACL entry. **Note**: This object applies only to extended ACLs.
snAgAclSourceOperand2	The second TCP or UDP port number in the source comparison, if applicable. If the value is 0, then this object does not apply to this ACL entry. **Note**: This object applies only to extended ACLs.
snAgAclDestinationIp	The destination IP address. **Note**: This object applies only to extended ACLs.
snAgAclDestinationMask	The destination IP network mask. **Note**: This object applies only to extended ACLs.
snAgAclDestinationOperator	The comparison operator, when the destination IP protocol is 6 (TCP) or 17 (UDP). The operator can be one of the following: • 0 – equal to • 1 – not equal to • 2 – less than • 3 – greater than • 4 – range • 7 – The operator is not defined **Note**: This object applies only to extended ACLs.
snAgAclDestinationOperand1	The first TCP or UDP port number in the destination comparison, if applicable. If the value is 0, then this object does not apply to this ACL entry. **Note**: This object applies only to extended ACLs.
snAgAclDestinationOperand2	The second TCP or UDP port number in the destination comparison, if applicable. If the value is 0, then this object does not apply to this ACL entry. **Note**: This object applies only to extended ACLs.

Table 3: snAgAclTable Objects (Continued)

Object	Description
snAgAclPrecedence	The IP precedence value, which can be one of the following: • 0 – routine • 1 – priority • 2 – immediate • 3 – flash • 4 – flashoverride • 5 – critical • 6 – internet • 7 – network • 8 – The IP precedence is not defined **Note**: This object applies only to extended ACLs.
snAgAclTos	The IP Type of Service (ToS) value, which can be one of the following: • 0 – normal • 1 – minMonetaryCost • 2 – maxReliability • 3 – tosValue3 • 4 – maxThroughput • 5 – tosValue5 • 6 – tosValue6 • 7 – tosValue7 • 8 – minDelay • 9 – tosValue9 • 10 – tosValue10 • 11 – tosValue11 • 12 – tosValue12 • 13 – tosValue13 • 14 – tosValue14 • 15 – tosValue15 • 16 – The IP ToS is not defined **Note**: This object applies only to extended ACLs.
snAgAclEstablished	For comparison of a TCP port, indicates whether the ACL also applies to TCP sessions that have already been established. The value can be one of the following: • 0 – This option is disabled or does not apply • 1 – This option is enabled

Table 3: snAgAclTable Objects (Continued)

Object	Description
snAgAclLogOption	Whether logging is enabled for packets that are denied by the ACL. The value can be one of the following: • 0 – This option is disabled or does not apply • 1 – This option is enabled
snAgAclStandardFlag	Indicates whether this is a standard ACL. • 0 – This is not a standard ACL • 1 – This is a standard ACL
snAgAclRowStatus	The configuration status of this row in the snAgAclTable. • 3 – Delete the entry • 4 – Add the entry

snAgAclBindToPortTable

This table contains the following objects for applying ACLs to interfaces.

Table 4: snAgAclBindToPortTable Objects

Object	Description
snAgAclPortNum	The interface number to which you want to apply (bind) the ACL.
snAgAclPortBindDirection	The traffic direction to which you want to apply the ACL. • 0 – outbound traffic • 1 – inbound traffic
snAgAclNum	The numerical ID of the ACL you want to apply.
snAgAclNameString	The name of the ACL you want to apply, if you used a name instead of a numerical ID to identify the ACL.
snAgBindPortListInVirtualInterface	The physical ports that are members of the virtual interface, if you are applying the ACL to a virtual interface.
snAgAclPortRowStatus	The configuration status of this row in the snAgAclBindToPortTable. • 3 – Delete the entry • 4 – Add the entry

New MIB Objects for CPU Utilization Statistics

The HP MIB contains the following new objects, which provide CPU utilization statistics. To use these objects, you need to use HP MIB file MIB07122.mib or later.

Table 5: CPU Utilization Statistics Objects

Object	Description
snAgGblCpuUtil1SecAvg	The average CPU utilization during the latest one-second interval.
snAgGblCpuUtil5SecAvg	The average CPU utilization during the latest five-second interval.
snAgGblCpuUtil1MinAvg	The average CPU utilization during the latest one-minute interval.

Higher Maximum Number of VLANs in a Single Spanning Tree

Software release 07.1.22 allows you to configure all the port-based VLANs on a device to use a single instance of STP. Previous releases allow up to 128 VLANs to use the same instance of STP.

USING THE CLI

To determine the maximum number of VLANs you can configure on a device, enter the **show default** values command. The “vlan” row in the “System Parameters” section shows the default maximum number of VLANs you can configure, the maximum currently allowed by the device, and maximum you can configure the device to allow. Here is an example:

```
HP9308# show default values
...<some lines omitted>
System Parameters     Default     Maximum     Current
...<some lines omitted>
vlan                  16          2048        16
```

To increase the maximum number of VLANs, enter a command such as the following at the global CONFIG level of the CLI:

```
HP9308(config)# system-max vlan 2048
```

Syntax: [no] system-max vlan <num>

The <num> parameter indicates the maximum number of VLANs and can be from 4 to the number listed in the Maximum column of the **show default values** display. You must enter a value that is divisible by 4. For example, 88, 132, and 136 are valid, but 35, 127, and 130 are not valid.

NOTE: Implementation of this enhancement includes moving the **spanning-tree single** command in the startup-config file to come before the VLAN configuration commands. In software releases earlier than 07.1.22, the **spanning-tree single** command comes after the VLAN configuration commands. If you boot the device using a software release earlier than 07.1.22, but with a startup-config file saved while running 07.1.22 or later, the CLI disregards the **spanning-tree single** command due to its unexpected position in the startup-config file and thus does not enable the feature.

Specifying Different Servers for Individual AAA Functions

In a RADIUS or TACACS+ configuration, you can designate a server to handle a specific AAA task. For example, you can designate one TACACS+ server to handle authorization and another TACACS+ server to handle accounting. For RADIUS, you can specify individual servers for authentication and accounting, but not for authorization. You can set the RADIUS or TACACS+ key for each server.

USING THE CLI

To specify different TACACS+ servers for authentication, authorization, and accounting:

```
HP9308(config)# tacacs-server host 1.2.3.4 auth-port 49 authentication-only key abc
HP9308(config)# tacacs-server host 1.2.3.5 auth-port 49 authorization-only key def
```

```
HP9308(config)# tacacs-server host 1.2.3.6 auth-port 49 accounting-only key ghi
```

Syntax: tacacs-server host <ip-addr> | <server-name> [authentication-only | authorization-only | accounting-only | default] [key <string>]

Syntax: radius-server host <ip-addr> | <server-name> [authentication-only | accounting-only | default] [key 0 | 1 <string>]

The **default** parameter causes the server to be used for all AAA functions.

USING THE WEB MANAGEMENT INTERFACE

To specify different TACACS+ servers for authentication, authorization, and accounting 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. Select the TACACS or RADIUS link. For TACACS+, the following panel is displayed:

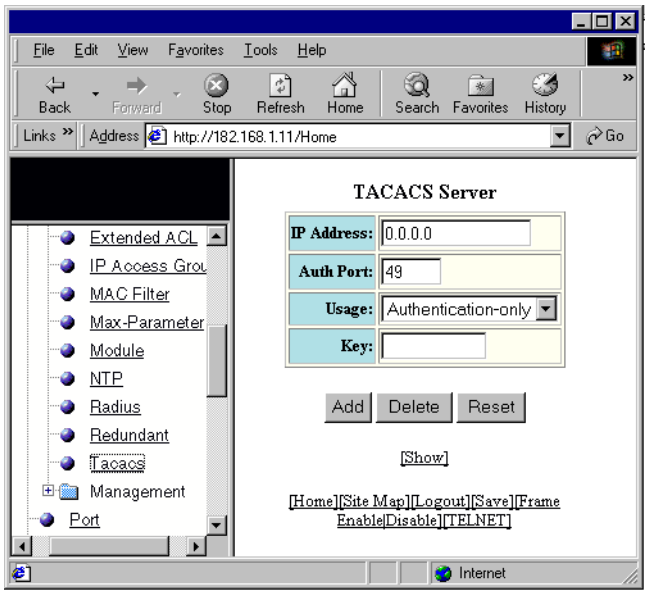

3. In the IP Address field, enter the IP address of the TACACS+ server.
4. In the Usage field, select Authentication-only, Authorization-only, Accounting-only, or All AAA functions.
5. Click Add to save the changes to the device's running-config.
6. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Encrypting RADIUS and TACACS+ Keys

When you display the configuration of the HP device, the RADIUS and TACACS keys are encrypted.

For example:

```
HP9308(config)# tacacs key 1 abc

HP9308(config)# write terminal
...
tacacs-server host 1.2.3.5 auth-port 49 key $!2d
```

Syntax: tacacs key 0 | 1 <string>

Syntax: radius key 0 | 1 <string>

Encryption of the RADIUS or TACACS keys is done by default. To disable encryption, use the **0** parameter.

Entering Privileged EXEC Mode After a Telnet or SSH Login

By default, you enter User EXEC mode after a successful login through Telnet or SSH. Optionally, you can configure the device so that you enter Privileged EXEC mode after a Telnet or SSH login. To do this, use the following command:

```
HP9308(config)# aaa authentication login privilege-mode
```

Syntax: aaa authentication login privilege-mode

Telnet/SSH Login Prompt Obtained from TACACS+ Server

When TACACS+ authentication is configured, the HP device now obtains both the login prompt and the password prompt from the TACACS+ server. Previously it obtained only the password prompt from the TACACS+ server.

Telnet/SSH Prompts When TACACS+ Server is Unavailable

When TACACS+ is the first method in the authentication method list, the device displays the login prompt received from the TACACS+ server. If a user attempts to login through Telnet or SSH, but none of the configured TACACS+ servers are available, the following takes place:

- If the next method in the authentication method list is "enable", the login prompt is skipped, and the user is prompted for the Enable password (that is, the password configured with the **enable super-user-password** command).
- If the next method in the authentication method list is "line", the login prompt is skipped, and the user is prompted for the Line password (that is, the password configured with the **enable telnet password** command).

Command Authorization and Accounting for Console Commands

The device now supports command authorization and command accounting for CLI commands entered at the console. To configure the device to perform command authorization and command accounting for console commands, enter the following:

```
HP9308(config)# enable aaa console
```

Syntax: enable aaa console

NOTE: If you have previously configured the device to perform command authorization using a RADIUS server, entering the **enable aaa console** command prevents the execution of any subsequent commands entered on the console. This happens because RADIUS command authorization requires a list of allowable commands from the RADIUS server. This list is obtained during RADIUS authentication. Since authentication is not performed on the console, this list is never obtained from the RADIUS server. Consequently, there are no allowable commands on the console.

Higher Maximum Number of IP Routes

Software release 07.1.22 increases the maximum number of IP routes that a redundant management module can have from 200,000 to 256,000. This number of routes is supported on Management II or higher modules with at least 128MB memory.

The default maximum number of IP routes is still 128,000. To increase the maximum number of routes the device can have, use the following CLI method.

USING THE CLI

To increase the maximum number of IP routes the device can have, enter a command such as the following:

```
HP9308(config)# system-max ip-route 256000
```

Syntax: system-max ip-route <num>

The <num> parameter specifies the maximum number of IP routes and can be from 4096 – 256000. The default for redundant management modules with at least 128MB memory is 128000.

Description of Enhancements in 07.1.19

Mini-GBIC Modules for the 9304M and 9308M

Software release 07.1.19 supports the J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module and the J4856A HP Procurve 9300 Mini-GBIC Module, two new eight-port forwarding modules for the 9304M and 9308M chassis. Both modules contain eight slots for miniature Gigabit Interface Converters (mini-GBICs). The modules provide flexibility by allowing you to install any combination of the following mini-GBICs:

- J4858A HP Procurve Gigabit-SX Mini-GBIC: 1000BaseSX port with an LC connector; supports multimode fiber cabling. Connection to a 1000BaseLX port is not supported.
- J4859A HP Procurve Gigabit-LX Mini-GBIC: 1000BaseLX port with an LC connector; supports single-mode and multi-mode fiber cabling. Connection to a 1000BaseSX port is not supported.

For more on this topic, refer to *Installing and Removing a Mini-GBIC*, provided with the mini-GBIC module and also available in the technical **support** | **manuals** area of the HP Procurve website at **http://www.hp.com/go/hpprocurve**.

Optimizing the Forwarding Cache for the Default Route

Software release 07.1.19 allows you to enhance routing performance in environments such as ISPs where the routing switch has a very large number of routes that use the default route to reach the next hop.

NOTE: This enhancement applies only to the 9308M/9304M routing switches WITH Redundant Management.

The 9304M and 9308M routing switches use Content Addressable Memory (CAM) as a fast lookup cache to optimize IP forwarding. The CAM contains an IP route's destination and the IP address of the next-hop gateway, as well as pointers to packet information in various system buffers. When the routing switch is ready to forward a packet to its destination, the routing switch checks the CAM for a forwarding entry for the packet.

- If the CAM contains an entry, the routing switch uses the entry to forward the packet.
- If the CAM does not contain an entry, the routing switch searches the IP route table for a route to the packet's destination, then programs an entry into the CAM for the destination and its next-hop gateway. The routing switch uses the CAM entry to forward the next packet to this destination.

By default, the CAM is optimized for environments with a lot of routes to different destination networks. Each CAM entry provides fast-path information for a different destination sub-net.

Software 07.1.19 provides a new option that optimizes forwarding in environments where the routing switch uses the default route to forward traffic for many of the destination networks. When you use the new option, the routing switch divides the default route into 4096 separate sub-nets, one for each entry in the CAM. Each of the entries has a 12-bit prefix. Routes that contain the same values for the first 12 bits of the address are aggregated together.

Entries that contain a sub-net that does not use the default route are "ineligible" to use one of the default-route CAM entries for forwarding. The routing switch uses the IP forwarding cache or the CPU for forwarding.

The default route can have more than one next-hop gateway address. When this is the case, the routing switch load balances traffic across the gateways using the IP load sharing settings in effect in the software. For information, see the "Configuring IP Load Sharing" section in the "Configuring IP" chapter of the *Book 1: Installation and Getting Started Guide*.

Enabling Default Route Optimization

To enable the feature, use the following CLI method.

USING THE CLI

To enable default route optimization, enter the following command at the global CONFIG level of the CLI:

```
HP9308(config)# ip net-aggregate
```

Syntax: [no] ip net-aggregate

Displaying CAM Entries

To display the entries in the CAM, use the following CLI method.

USING THE CLI

To display the entries in the CAM, enter the following command at any level of the CLI:

```
HP9308(config)# show ip net-aggregate
Total prefixes: 4096, CAM Ineligible: 31, Setups: 14, Updates 7477
Start index: 1
0.0.0.0/12        Gateway: 101.77.7.101
CAM Entry Flag: 00000703H
CIDX0: 5553  CIDX8: 4127
0.16.0.0/12        Gateway: 101.78.7.101
CAM Entry Flag: 00000003H
CIDX0: 5552
0.32.0.0/12        Gateway: 101.79.7.101
CAM Entry Flag: 00000003H
CIDX0: 5551
0.48.0.0/12        Gateway: 101.76.7.101
CAM Entry Flag: 00000003H
CIDX0: 5550
0.64.0.0/12        Gateway: 101.77.7.101
CAM Entry Flag: 00000003H
CIDX0: 5549
0.80.0.0/12        Gateway: 101.78.7.101
CAM Entry Flag: 00000003H
CIDX0: 5548
0.96.0.0/12        Gateway: 101.79.7.101
CAM Entry Flag: 00000003H
CIDX0: 5547
0.112.0.0/12        Gateway: 101.76.7.101
CAM Entry Flag: 00000003H
CIDX0: 5546
--More--, next page: Space, next line: Return key, quit: Control-c
```

As shown by this example, the default-route optimization feature divides the default route into individual networks with 12-bit prefixes. The first entry is network 0.0.0.0/12, the second entry is network 0.16.0.0/12, and so on. If a destination network is reachable through the default route, the routing switch uses the corresponding CAM entry to forward traffic to the destination. If the destination is reachable through a route other than the default route, the destination is ineligible to use the CAM for fast-path forwarding. The routing switch uses the IP forwarding cache or the CPU to forward the traffic instead.

Syntax: show ip net-aggregate [<starting-entry-num> | <ip-addr> | not-eligible]

The <starting-entry-num> specifies the entry number you want the command's output to start with. By default, the display begins with the first entry.

The <ip-addr> parameter specifies the IP address of a destination. The CAM entry that contains the specified address is displayed.

The **not-eligible** parameter displays only the entries that are ineligible for use because they contain a destination network that the routing switch uses a route other than the default route to reach. To display forwarding information for ineligible entries, use the following commands:

- **show ip cache** – see the "Displaying the Forwarding Cache" section "Configuring IP" chapter of the *Book 1: Installation and Getting Started Guide*.
- **show ip route** – see the "Displaying the IP Route Table" section "Configuring IP" chapter of the *Book 1: Installation and Getting Started Guide*.

The **show ip net-aggregate** command displays the following information.

Table 6: CLI Display of CAM

This Field...	Displays...
Total prefixes	The total number of entries in the CAM.
CAM Ineligible	The number of entries that cannot be used for fast-path forwarding because the IP route table contains a route whose destination network is contained in the entry's aggregate network, but does not use the default route.
Setups	The number of times the entire CAM has been reprogrammed during the current power cycle. Generally, this occurs when the default route changes.
Updates	The number of individual entries that have been updated due during the current power cycle to a route change.
Start index	The entry number of the first entry in the display. If you specify a starting entry number when you enter the **show ip net-aggregate** command, then this field shows that number. Otherwise, the starting number is 1.
Destination address	An aggregate network address. If a route's destination is contained in this aggregate address, then this CAM entry is applicable to the destination. **Note**: When default-route optimization is enabled, the entry is actually used only if the routing switch uses the default route to reach the destination.
Gateway	The IP address of the next-hop gateway reached through the default route. **Note**: The default route can have more than one next-hop gateway address. When this is the case, the routing switch load balances traffic across the gateways using the IP load sharing settings in effect in the software. For information, see the "Configuring IP Load Sharing" section in the "Configuring IP" chapter of the *Book 1: Installation and Getting Started Guide*.
CAM Entry Flag	A value used by HP Technical Support for troubleshooting.
CIDX*n*	A value used by HP Technical Support for troubleshooting.

BGP4 Next-Hop Recursion

For each BGP4 route a routing switch learns, the routing switch performs a route lookup to obtain the IP address of the route's next hop. A BGP4 route becomes eligible for installation into the IP route table only if the following conditions are true:

- The lookup succeeds in obtaining a valid next-hop IP address for the route.
- The path to the next-hop IP address is an Interior Gateway Protocol (IGP) path or a static route path.

By default, the software performs only one lookup for a BGP route's next-hop IP address. If the next-hop lookup does not result in a valid next-hop IP address or the path to the next-hop IP address is a BGP path, the software considers the BGP route's destination to be unreachable. The route is not eligible to be installed in the IP route table.

It is possible for the BGP route table to contain a route whose next-hop IP address is not reachable through an IGP route, even though a hop farther away can be reached by the routing switch through an IGP route. This can occur when the IGPs do not learn a complete set of IGP routes, resulting in the routing switch learning about an internal route through IBGP instead of through an IGP. In this case, the IP route table does not contain a route that can be used to reach the BGP route's destination.

To enable the routing switch to find the IGP route to a BGP route's next-hop gateway, enable recursive next-hop lookups. When you enable recursive next-hop lookup, if the first lookup for a BGP route results in an IBGP path originated within the same Autonomous System (AS), rather than an IGP path or static route path, the routing switch performs a lookup on the next-hop gateway's next-hop IP address. If this second lookup results in an IGP path, the software considers the BGP route to be valid and thus eligible for installation in the IP route table. Otherwise, the routing switch performs a lookup on the next-hop IP address of the next-hop gateway's next hop, and so on, until one of the lookups results in an IGP route.

Example When Recursive Route Lookups Are Disabled

Here is an example of the results of an unsuccessful next-hop lookup for a BGP route. In this case, next-hop recursive lookups are disabled. The example is for the BGP route to network 240.0.0.0/24.

```
HP9308# show ip bgp route
Total number of BGP Routes: 5
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
       Network          ML Next Hop        Metric     LocPrf      Weight Status
1      0.0.0.0          0  10.1.0.2        0          100         0      bI
2      102.0.0.0        24 10.0.0.1        1          100         0      BI
3      104.0.0.0        24 10.1.0.2        0          100         0      BI
4      240.0.0.0        24 102.0.0.1       1          100         0      I
5      250.0.0.0        24 209.157.24.1    1          100         0      I
```

In this example, the routing switch cannot reach 240.0.0.0/24, because the next-hop IP address for the route is an IBGP route instead of an IGP route, and thus is considered unreachable by the routing switch. Here is the IP route table entry for the BGP route's next-hop gateway (102.0.0.1/24):

```
HP9308# show ip route 102.0.0.1
Total number of IP routes: 37
      Network Address   NetMask           Gateway          Port     Cost    Type
      102.0.0.0         255.255.255.0     10.0.0.1         1/1      1       B
```

The route to the next-hop gateway is a BGP route, not an IGP route, and thus cannot be used to reach 240.0.0.0/24. In this case, the routing switch tries to use the default route, if present, to reach the sub-net that contains the BGP route's next-hop gateway.

```
HP9308# show ip route 240.0.0.0/24
Total number of IP routes: 37
      Network Address   NetMask           Gateway          Port     Cost    Type
      0.0.0.0           0.0.0.0           10.0.0.202       1/1      1       S
```

Example When Recursive Route Lookups Are Enabled

When recursive next-hop lookups are enabled, the routing switch recursively looks up the next-hop gateways along the route until the routing switch finds an IGP route to the BGP route's destination. Here is an example.

```
HP9308# show ip bgp route
Total number of BGP Routes: 5
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
       Network          ML Next Hop        Metric     LocPrf      Weight Status
1      0.0.0.0          0  10.1.0.2        0          100         0      bI
2      102.0.0.0        24 10.0.0.1        1          100         0      BI
3      104.0.0.0        24 10.1.0.2        0          100         0      BI
4      240.0.0.0        24 102.0.0.1       1          100         0      BI
5      250.0.0.0        24 209.157.24.1    1          100         0       I
```

The first lookup results in an IBGP route, to network 102.0.0.0/24:

```
HP9308# show ip route 102.0.0.1
Total number of IP routes: 38
      Network Address   NetMask           Gateway           Port     Cost    Type
      102.0.0.0         255.255.255.0     10.0.0.1          1/1      1       B
```

Since the route to 102.0.0.1/24 is not an IGP route, the routing switch cannot reach the next hop through IP, and thus cannot use the BGP route. In this case, since recursive next-hop lookups are enabled, the routing switch next performs a lookup for 102.0.0.1's next-hop gateway, 10.0.0.1:

```
HP9308# show ip bgp route 102.0.0.0
Number of BGP Routes matching display condition : 1
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
     Network         ML Next Hop        Metric     LocPrf     Weight Status
1      102.0.0.0       24 10.0.0.1        1          100        0      BI
```

The next-hop IP address for 102.0.0.1 is not an IGP route, which means the BGP route's destination still cannot be reached through IP. The recursive next-hop lookup feature performs a lookup on 10.0.0.1's next-hop gateway:

```
HP9308# show ip route 10.0.0.1
Total number of IP routes: 38
      Network Address   NetMask           Gateway           Port     Cost    Type
      10.0.0.0          255.255.255.0     0.0.0.0           1/1      1       D
```

This lookup results in an IGP route. In fact, this route is a directly-connected route. As a result, the BGP route's destination is now reachable through IGP, which means the BGP route is eligible for installation in the IP route table. Here is the BGP route in the IP route table:

```
HP9308# show ip route 240.0.0.0/24
Total number of IP routes: 38
      Network Address    NetMask          Gateway           Port     Cost    Type
      240.0.0.0          255.255.255.0    10.0.0.1          1/1      1       B
```

This routing switch can use this route because the routing switch has an IP route to the next-hop gateway. Without recursive next-hop lookups, this route would not be in the IP route table.

Enabling Recursive Next-Hop Lookups

The recursive next-hop lookups feature is disabled by default. To enable the feature, use the following CLI method.

USING THE CLI

To enable recursive next-hop lookups, enter the following command at the BGP configuration level of the CLI:

```
HP9308(config-bgp-router)# next-hop-recursion
```

Syntax: [no] next-hop-recursion

SNMP MIB Objects for PIM SM

Software release 07.1.19 adds support for the following SNMP MIB objects:

- snPimJoinPruneInterval – the interval for join and prune messages
- snPimCandidateBSRTable – the candidate Bootstrap Router (BSR) table
- snPimRPSetTable – the Rendezvous Point (RP) set table
- snPimCandidateRPTable – the candidate RP table

These objects are customized for HP devices and are located in the HP MIB file, MIB071*xx*.mib, where *xx* is the software release the MIB file supports. To use these new PIM SM objects, you need MIB07118.mib or higher.

TACACS+ Exec Authorization Supports Non-HP A-V Pairs

To set a user's privilege level using a TACACS+ server, the HP device can accept either a HP-specific A-V pair or a non-HP-specific A-V pair. For more information, see "Configuring an Attribute-Value Pair on the TACACS+ Server", below.

Configuring an Attribute-Value Pair on the TACACS+ Server

During TACACS+ exec authorization, the HP device expects the TACACS+ server to send a response containing an A-V (Attribute-Value) pair that specifies the privilege level of the user. When the HP device receives the response, it extracts an A-V pair configured for the Exec service and uses it to determine the user's privilege level.

To set a user's privilege level, you can configure the "hp-privlvl" A-V pair for the Exec service on the TACACS+ server. For example:

```
user=bob {
  default service = permit
   member admin
   # Global password
  global = cleartext "cat"
   service = exec {
     hp-privlvl = 0
        }
}
```

In this example, the A-V pair `hp-privlvl = 0` grants the user full read-write access. The value in the hp-privlvl A-V pair is an integer that indicates the privilege level of the user. Possible values are 0 for super-user level, 4 for port-config level, or 5 for read-only level. If a value other than 0, 4, or 5 is specified in the hp-privlvl A-V pair, the default privilege level of 5 (read-only) is used. The hp-privlvl A-V pair can also be embedded in the group configuration for the user. See your TACACS+ documentation for the configuration syntax relevant to your server.

If the hp-privlvl A-V pair is not present, the HP device extracts the last A-V pair configured for the Exec service that has a numeric value. The HP device uses this A-V pair to determine the user's privilege level. For example:

```
user=bob {
   default service = permit
   member admin
   # Global password
   global = cleartext "cat"
   service = exec {
     privlvl = 15
        }
}
```

The attribute name in the A-V pair is not significant; the HP device uses the last one that has a numeric value. However, the HP device interprets the value for a non-"hp-privlvl" A-V pair differently than it does for a "hp-privlvl" A-V pair. The following table lists how the HP device associates a value from a non-"hp-privlvl" A-V pair with a HP privilege level.

Table 7: HP Equivalents for non-"hp-privlvl" A-V Pair Values

Value for non-"hp-privlvl" A-V Pair	hp Privilege Level
15	0 (super-user)
From 14 – 1	4 (port-config)
Any other number or 0	5 (read-only)

In the example above, the A-V pair configured for the Exec service is `privlvl = 15`. The HP device uses the value in this A-V pair to set the user's privilege level to 0 (super-user), granting the user full read-write access.

In a configuration that has both a “hp-privlvl” A-V pair and a non-”hp-privlvl” A-V pair for the Exec service, the non-”hp-privlvl” A-V pair is ignored. For example:

```
user=bob {
   default service = permit
   member admin
   # Global password
   global = cleartext "cat"
   service = exec {
     hp-privlvl = 4
     privlvl = 15
        }
}
```

In this example, the user would be granted a privilege level of 4 (port-config level). The `privlvl = 15` A-V pair is ignored by the HP device.

If the TACACS+ server has no A-V pair configured for the Exec service, the default privilege level of 5 (read-only) is used.

Encryption of BGP4 MD5 Authentication Keys

This enhancement is not described elsewhere in the product documentation. For descriptions of other enhancements, refer to the manuals provided for software release 06.6.*xx* and 07.1.*xx*. (See page 9.)

When you configure a BGP4 neighbor or neighbor peer group, you can specify an MD5 authentication string for authenticating packets exchanged with the neighbor or peer group of neighbors.

In software releases earlier than 07.1.19, the MD5 authentication string is displayed as clear text in the output of the following commands:

- **show running-config** (or **write terminal**)
- **show configuration**
- **show ip bgp config**

For added security, software release 07.1.19 encrypts display of the authentication string by default. The software also provides an optional parameter to disable encryption of the authentication string, on an individual neighbor or peer group basis.

When encryption of the authentication string is enabled, the string is encrypted in the CLI regardless of the access level you are using.

If you are upgrading a device that is already configured for BGP4, when you save the configuration to the startup-config file, the software automatically converts the command syntax for BGP4 neighbors and peer groups into the new syntax that includes the encryption option. If you display the running-config after reloading with software release 07.1.19, the BGP4 commands that specify an authentication string show the string in encrypted form.

In addition, when you save the configuration to the startup-config file, the file contains the new BGP4 command syntax and encrypted passwords or strings.

CAUTION:HP recommends that you save a copy of the startup-config file for each routing switch you plan to upgrade. If you need to return to a software release earlier than 07.1.19, the earlier software will not recognize the passwords or authentication keys in their encrypted form and will not be able to convert them back to their clear form.

Encryption Example

The following commands configure a BGP4 neighbor and a peer group, and specify MD5 authentication strings (passwords) for authenticating packets exchanged with the neighbor or peer group.

```
HP9300(config-bgp-router)# local-as 2
HP9300(config-bgp-router)# neighbor xyz peer-group
HP9300(config-bgp-router)# neighbor xyz password abc
HP9300(config-bgp-router)# neighbor 10.10.200.102 peer-group xyz
HP9300(config-bgp-router)# neighbor 10.10.200.102 password test
```

Here is how the commands appear when you display the BGP4 configuration commands:

```
HP9300(config-bgp-router)# show ip bgp config
Current BGP configuration:
router bgp
 local-as 2
 neighbor xyz peer-group
 neighbor xyz password 1 $!2d
 neighbor 10.10.200.102 peer-group xyz
 neighbor 10.10.200.102 remote-as 1
 neighbor 10.10.200.102 password 1 $on-o
```

Notice that the software has converted the commands that specify an authentication string into the new syntax (described below), and has encrypted display of the authentication strings.

New Command Syntax

Since the default behavior in software release 07.1.19 does not affect the BGP4 configuration itself but does encrypt display of the authentication string, the CLI does not list the encryption options.

Syntax: [no] neighbor <ip-addr> | <peer-group-name> password [0 | 1] <string>

The <ip-addr> | <peer-group-name> parameter indicates whether you are configuring an individual neighbor or a peer group. If you specify a neighbor's IP address, you are configuring that individual neighbor. If you specify a peer group name, you are configuring a peer group.

The **password** <string> parameter specifies an MD5 authentication string for securing sessions between the routing switch and the neighbor. You can enter a string up to 80 characters long. The string can contain any alphanumeric characters, but the first character cannot be a number. If the password contains a number, do not enter a space following the number.

The **0 | 1** parameter is the encryption option, which you can omit (the default) or which can be one of the following.

- **0** – Disables encryption for the authentication string you specify with the command. The password or string is shown as clear text in the output of commands that display neighbor or peer group configuration information.
- **1** – Assumes that the authentication string you enter is the encrypted form, and decrypts the value before using it.

NOTE: If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter 0 or 1. Instead, omit the encryption option and allow the software to use the default behavior.

If you specify encryption option 1, the software assumes that you are entering the encrypted form of the password or authentication string. In this case, the software decrypts the password or string you enter before using the value for authentication. If you accidentally enter option 1 followed by the clear-text version of the password or string, authentication will fail because the value used by the software will not match the value you intended to use.

NOTE: For simplicity, the other optional parameters with the **neighbor** command are not shown. For complete syntax information, see the "Configuring BGP4" chapter in the *Installation and Getting Started Guide* or the *Command Line Interface Reference*.

Displaying the Authentication String

If you want to display the authentication string, enter the following commands:

```
HP9300(config)# enable password-display
HP9300(config)# show ip bgp neighbors
```

The **enable password-display** command enables display of the authentication string, but only in the output of the show ip bgp neighbors command. Display of the string is still encrypted in the startup-config and running-config files. Enter the command at the global CONFIG level of the CLI.

NOTE: The command also displays SNMP community strings in clear text, in the output of the **show snmp server** command.

Software Fixes

This section lists the problems that have been fixed in software releases 07.1.10 – 07.1.24. For information about fixes in an earlier software release, see the release notes for that release.

NOTE: Software releases sometimes apply only to specific products. Check the list of products supported by the release to make sure the release applies to your product. Each set of release notes lists the products to which the release applies.

Fixed in 07.1.24

- **AAA** – Named access lists with network mask IDs were not written to the running-config properly when AAA was enabled.
- **AAA** – During Telnet login, the message “AAA operation in progress..please try again after some time” was constantly displayed.
- **AAA** – When AAA command authorization or accounting was configured, you could not delete IP prefix lists from the running-config.
- **AAA** – With AAA enabled, pasting a large configuration file from Telnet caused pasting to loop continuously.
- **ACL** – A software issue could cause session entries for ACLs to quickly be consumed, after which ACL traffic was replicated and sent to the CPU, causing high CPU utilization.
- **ACLs** – If you configured an ACL whose address range was contained within the address range configured in another ACL, the software did not filter the traffic based on the new ACL.
- **ACLs** – ACLs dropped fragmented TCP or UDP packets.
- **ACLs** – Packets denied by an ACL caused “Authentication Failure” syslog messages and traps to be generated, in addition to “ACL Deny” syslog messages and traps. Now only “ACL Deny” syslog messages and traps are generated.
- **ACLs** – The device did not forward fragmented packets that were out of sequence. This issue affected ACLs applied to inbound traffic or outbound traffic. This occurred for the following reasons:
  - The outbound ACL code was missing the fragmentation packet check.
  - For inbound ACLs, the code took a fast path that did not create an ACL session. Without the ACL session, fragmented packets were dropped. Due to possible fragmentation, the code did not take the fast path.
  - The ACL engine incorrectly identified the first fragmented packet (more_bit on and offset 0) to be a fragment. Hence, this packet was dropped and no session was created.
- **ACLs** – Performance was affected when ACLs were applied to interfaces.
- **AppleTalk** – The device could not learn more than 300 AppleTalk routes, even if the **system-max atalk-**

routes <num> command was used to allow the maximum number of routes.

- **ARP** – A gratuitous ARP was not sent out the primary interface when it was configured with multiple IP addresses. This issue affected Gigabit ports only.
- **BGP4** – In certain cases, if BGP4 next-hop recursion ended in a recursion loop, periodic spikes in CPU utilization could occur. The loops could occur as the result of a next-hop dependency loop of some routes; for example, BGP4 routes 209.1.44.210/32 (next-hop 209.1.44.211) and 209.1.44.211/32 (next-hop 209.1.44.210).
- **CLI** – In some cases, entering the **show tech** command from a Telnet session, then entering any other show command that uses a paged display and ending the display using CTRL+C could cause system errors, such as hanging of the management session. The errors could occur because the **show tech** command uses a TCP buffer timer to ensure that it has enough buffers for sending the command's output to the Telnet session. If another subsystem in the software had started using the same memory location as the timer when you entered Ctrl+c, Ctrl+c cleared that memory location and thus could cause errors in the subsystem that was using the memory location.

NOTE: The same type of issue, where one subsystem had started a timer but the timer was still in use when another subsystem inadvertently canceled the timer, could cause errors in other subsystems. The cause of the errors is fixed in 07.1.24.

- **CLI** – If command authorization and/or accounting was configured, you could not paste commands into the device's running configuration.
- **CLI** – If the show ip ospf config command was entered in more than one management session (serial or Telnet) at the same time, the software could reload. This occurred after a couple of attempts to simultaneously enter the command.
- **CLI** – If you entered the show statistics ethernet 1/1 command, the InPkts and OutPkts counters were displayed in 32-bit values, even though the internal counters were stored in 64-bit values.
- **CLI** – If you accidentally entered the **ping** command followed by a string that began with a dot (for example **ping .hfsdjhf**), the device did not send a ping packet, but it did consume enough buffers to construct a packet and did not free the buffers. If the **ping** command was entered about 20 more times with a value beginning with a dot, all the buffers were consumed, preventing successful transmission of a ping packet when a valid IP address was entered with the **ping** command.
- **CLI** – The **uplink-switch ethernet** command did not appear in the running-config or startup-config file, even after you saved the configuration. The configuration was present in the active configuration but disappeared from the configuration as soon as you exited the CONFIG mode. However, the uplink ports were listed correctly by the **show vlan** command.
- **CLI** (Mini-GBIC module only)– The current chassis temperature displayed by the **show chassis** command was inaccurate. Note that the temperature warning and shutdown functions use the actual temperature of the chassis to determine whether the chassis temperature has crossed the warning or shutdown thresholds, not the temperature value reported by the **show chassis** command.
- **CLI** – If the startup-config file contained the command **ip ssh source-interface**, the following error message as displayed following a reload "Invalid tcb access 0x00000000 at 20135098 20452bd0 204509d4".
- **CLI** – Named ACLs did not appear in the running-config following a software reload.
- **CLI** – Named ACLs were not properly written to the running-config if AAA was enabled.
- **CLI** – The **access-list** command was not listed in response to entering **?** at the CLI prompt at the global CONFIG level, even though the command is valid.
- **CLI** – The **no router ospf** command could cause memory corruption.
- **CLI** – If the network mask format was changed from CIDR to classical sub-net mask format, the CLI changed ACL definitions, resulting in incorrect addresses in the ACLs.
- **DHCP** – DHCP packets longer than 1500 bytes caused loss of one DRAM buffer per such packet. Once the device ran out of buffers, the software reloaded.

- **DHCP** – The HP device forwarded a DHCP offer packet to a client using a MAC address in the ARP table instead of using the address in the chaddr field of the offer packet. In some cases, this could cause the HP device to send a DHCP offer to the wrong client. This occurred when more than one IP helper address was configured on the HP interface, and the device received offers for the same client from more than one DHCP server.

  In this case, the HP device created a separate ARP entry for each offer. The multiple ARP entries associated the same MAC address with all of the offered IP addresses. When a second client sent a DHCP request, if one of the DHCP servers offered an IP address that was offered to the first client but not accepted, the HP device used the ARP table to forward the offer. If the ARP table still had an entry that associated the offered address with the first client, the HP device forwarded the offer to the first client, even though the offer should have been forwarded to the second client.

  This issue occurred only in configurations where the DHCP server was multiple hops away from the HP device, and only before the ARP entry for the second client aged out.
- **DHCP** – The routing switch dropped a DHCP NAK packet sent by an NT DHCP server instead of forwarding the packet to the client.

  NOTE: This problem occurred only in software release 07.1.22.
- **Fast Ethernet statistics** – The counters for 100 Mbps interface statistics were 32-bit counters, which resulted in rapid wrapping of the values for higher-speed interfaces. The counters are now 64-bit counters. The corresponding SNMP MIB objects for these counters are the following: ifHCInOctets, ifHCInUcastPkts, ifHCInMulticastPkts, ifHCInBroadcastPkts, ifHCOutOctets, ifHCOutUcastPkts, ifHCOutMulticastPkts, ifHCOutBroadcastPkts.
- **IP** – When the default route is deleted, there sometimes is an interval during which there is no default route before the new default route gets learned. During this interval, the device drops packets for flows for which there is no route. Although this is the correct behavior, in the previous software release the device continued to drop packets for the flows even after a new default route was installed.
- **IP** – If the device had both a RIP default route and an IP static route that was the default route and the administrative distance of the RIP default route was lower, when the RIP default route was removed or was no longer available, the device did not restore the backup default route (the static IP route). This issue occurred only once following a software reload; entering the clear ip route command would restore the static default route. Afterwards, the issue did not happen again.
- **IP and OSPF** – A bug in packet reception data caching caused two bytes at the end of packets to be corrupted. When this occurred, the message "OSPF: intf rcvd bad pkt: ..." was sent to the Syslog.
- **IP multicast** – If a receiver's VLAN did not contain a tagged port and multicast traffic was received from a tagged port, the multicast packet could be corrupted with an invalid IP header checksum.
- **IP multicast** – If IP multicast routing was enabled on VLANs that shared a tagged port, and receivers were on the same VLAN as the sender and also on the other VLAN, received packets could be corrupted with 16 extra bytes of data, prepended to the original packet. This issue affected both PIM and DVMRP.
- **IP Multicast routing** – When PIM DM multicast was running with virtual interfaces on an untagged VLAN port, if an IGMP join message was received on one of the virtual interfaces, packets were copied to the CPU even if there were no local clients on the virtual interface sending the multicast data. This could result in performance degradation. This fix improves the performance of multicast forwarding when using virtual interfaces on an untagged port.
- **IPX** – The device could reset when deleting a cache entry.
- **IPX** – If you disabled a port that had directly-connected IPX networks, entered the **clear ipx route** command and then re-enabled the port, the software created two IPX route table entries for the route, one with a hop count of 0 and the other with a hop count of 16. The **clear ipx route** command could not delete the phantom route with the hop count of 16.
- **IPX** – Traffic was not forwarded over a Layer 2 port-based VLAN if an IPX protocol VLAN was configured with virtual interfaces.

- **IPX** – Packet buffers were freed twice for forwarded IPX packets.
- **OSPF** – If a static route was configured with a non-zero host portion and then was redistributed into OSPF, the route could not later be removed by the **no redistribute static** command.
- **OSPF** – If an OSPF interface that was down was configured with a priority of 0 and was set to the passive mode afterwards, the HP device incorrectly advertised the attached network as a reachable route in OSPF.
- **OSPF** – In some cases, a learned route could be preferred over a directly connected route. In the following example, R3 could incorrectly select a path through R2 or R4 even though the path through R3 was the directly connected path.

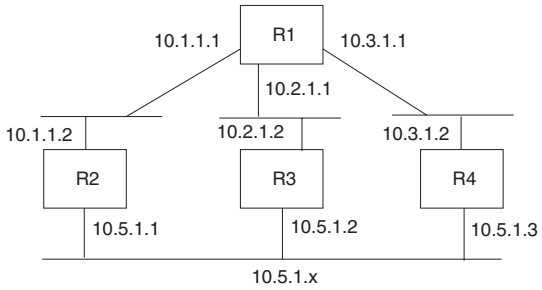

- **OSPF** – If both OSPF and BGP4 were advertising a default route from another router with multiple paths, removing OSPF default route advertisement did not restore a default route with a BGP4 path if ports to the path had flapped.
- **OSPF** – In a configuration where a routing switch was configured as an Area Border Router (ABR) between a Not-So-Stubby Area (NSSA) and a backbone area, the routing switch did not pass the default route from the NSSA ASBR (the router with default-information-originate enabled) to the backbone area.
- **OSPF** – The **default-information-originate** command's **metric-type** option did not work.
- **PIM Dense** – When IP multicast was running with PIM Dense on VLANs that had more than one tagged port, and there was more than one physical path from the multicast stream sender to the HP device on the same virtual interface, if the incoming physical interface of the multicast stream changed to another one but was still in the same virtual interface, receivers on downstream interfaces could not receive data.
- **PIM Dense** – When IP multicast was running with PIM Dense using sub-net VLANs that overlapped other sub-net VLANs, clients connected to non-PIM DR routers couldn't join groups and receive traffic.
- **PIM Dense** – When IP multicast was running with PIM DM on VLANs that had more than one tagged port, if PIM Dense lost assert on one physical port, clients connected to other physical ports in the same virtual interface stopped receiving traffic.
- **PIM Dense** – When IP multicast was running on VLANs that shared tagged ports, if there was a client in the same VLAN as the server, clients that were on other VLANs but were connected to the same physical port received extra copies of the packet.
- **PIM SM** – If the routing switch had more than one equal-cost path to a network, and PIM SM was disabled on the interface on the selected path, unicast routing attempted to reach the multicast source, causing display of the error message "invalid vif". In software release 07.1.24, the software selects only the paths that are on interfaces on which PIM SM is enabled.
- **RIP** – If the device received an update for a RIP route where the next hop was different from the address of the advertising interface, the HP device installed the route in the route table with the advertising interface as the next hop. In the current release, the next hop is obtained from the RIP update packet.
- **Route health injection** – The feature did not work properly. When the ARP entry for a host timed out, the routing switch was not able to install the ARP reply immediately. This caused the injected route to be removed from the upstream router. Also, the **ip dont-advertise** command did not inject the network route into the upstream router, regardless of the result of the health check.
- **RIP** – A route was not updated if its originator was different from its next hop in RIP V2.
- **SNMP** – The MTU for 10/100 and Gigabit Ethernet interfaces was not reported correctly.

- **SNMP** – An SNMP walk for the object snPimCandidateBSRTable failed if PIM was not configured on the device.
- **SNMP** – The SNMP agent for tables snRtIpPortAddrTable, snRtIpRipPortConfigTable, and snRtIpPortConfigTable looped for virtual interface numbers higher than 256.
- **SNMP** – There was continuous looping of the SNMP agent for the tables snRtIpRipPortConfigTable, snRtIpPortConfigTable, and snCandidateBSRTable.
- **SSH** – The software now supports Perl-generated user names longer than 1,000 bytes.
- **SSH** – The software now frees resources when an error occurs with RSA key generation, as well as when an SSH session is terminated.
- **SSH** – The **crypto key generate rsa** command now works when entered from an SSH session.
- **SSH** – The **show ip pim mcache** <group-address> command could cause the SSH connection to be dropped if the group had more than around 25 cache entries.
- **STP** – If single STP was enabled on the device, but STP was disabled on port that was in a VLAN on which STP was enabled, the port would come up as a blocked STP port.
- **STP** – The device sent a TCN packet when the state changed on a port configured for Fast Port Span.
- **Syslog** – Port state changes (between up and down) for Gigabit ports in a trunk group were not reported to the Syslog correctly.
- **TACACS+** – When TACACS+ authentication was configured, if a user's Telnet session timed out, and the user entered an invalid user name and password, the user was still granted access to the device.
- **TACACS+** – Telnet authentication did not display prompts provided by the TACACS+ server, other than username and password.
- **TACACS+** – If the TACACS+ key was invalid, the software did not try the next server or the next authentication method, but instead displayed an error message. In software release 07.1.24, the software tries the next server or, if there is no other server, tries the next authentication method.
- **Telnet** – When lines of commands were pasted into the device's running configuration, the newline character was filtered incorrectly, causing the commands to be displayed incorrectly.
- **Telnet/SSH** – The software now checks if the Telnet/SSH TCB pointer is valid before passing it to the TCP code. If an illegal access attempt is detected, a message is displayed on the console and written to syslog.
- **Telnet** – When lines of commands were pasted into the device's running configuration, the newline character was filtered incorrectly, causing the commands to be displayed incorrectly.
- **Telnet/SSH** – The software now checks if the Telnet/SSH TCB pointer is valid before passing it to the TCP code. If an illegal access attempt is detected, a message is displayed on the console and written to syslog.
- **Traceroute** – The **traceroute** command did not list routers outside the local network if the network was connected to the Internet by a NAT device. This issue occurred only with some NAT devices, which made non-standard changes to ICMP time exceeded packets. The HP software was enhanced to work around the non-standard packet modifications.
- **Traceroute** – If you specified a hostname instead of an IP address with the **traceroute** command, the software reloaded.
- **Trunk groups** – In a four-port trunk group, a MAC address learned on the trunk sometimes was aged out improperly, even when traffic was actively using the MAC address. A side effect of this issue was that occasionally, a small percentage of packets were dropped when there was heavy traffic going into the trunk group.
- **Trunk groups** – In test situations where traffic was continuously sent to trunk ports at line rate, normal MAC aging and relearning could cause packet loss while the CPU was relearning aged out MAC entries.
- **Trunk groups** – If VRRP was enabled on a Gigabit trunk group, Layer 3 traffic from the device that was using the HP device as a gateway was forwarded by the CPU instead of by the hardware.
- **VLANs** – In a configuration where separate hosts connected to two port-based VLANs each had the same IP

address, pings from a host to the virtual routing interface on its VLAN stopped after awhile. This occurred because the routing switch created a MAC entry and an ARP entry for the same host in the other VLAN.

- **VLANs and STP** – A tagged port leaked unicast traffic even if the port was blocked.
- **VRRP** – Trying to use the **no backup** command when configuring a VRRP or VRRPE VRID on an interface before the VRID itself was configured caused the software to reload.
- **VRRP** – If VRRP was enabled and you hot-swapped a forwarding module on which multiple VRIDs on virtual interfaces that spanned multiple modules were configured, the hosts using the VRIDs could no longer ping one another.
- **VRRP** – A disconnected or disabled owner VRRP interface would send out two gratuitous ARPs when it was connected or re-enabled, one for the physical interface and one for the virtual interface, both with the same IP address. This would cause the backup VRRP router to respond to the ARP for the physical interface (if still active), which could cause come inconsistency in the devices receiving these ARP requests and replies.
- **Web management interface** – If you set access restrictions on the Web management interface, these restrictions were not reflected until you saved the configuration and reloaded it.

Fixed in 07.1.22

- **BGP4** – When comparing routes, the software interpreted the path length as the number of AS entries in the AS-SET. In software release 07.1.22, the software interprets the path length of the AS-SET as 1.
- **BGP4** – If a route map used by a redistribution command was deleted, the software continued to advertise the routes that matched the route map. In software release 07.1.22, if you delete a route map, you can cause the software to stop redistributing the routes that match the deleted route map, by entering the **clear ip bgp local route** command. See “Refreshing Routes Redistributed into BGP4” on page 33.
- **BGP4** – The first time BGP4 was configured on the device, you needed to configure the local AS, then enter the **clear ip bgp neighbor all** command, or save the configuration changes and reload the software. Otherwise, BGP4 would not advertise the locally originated routes to IBGP neighbors. If any BGP4 neighbors were configured, the **clear ip bgp neighbor all** command would reset all the neighbor sessions.
- **BGP4** – In a configuration where the routing switch used a password to validate a BGP4 session with a neighbor, if the neighbor rebooted while the session was in the Established state, the neighbor could cause a BGP4 error on the routing switch after rebooting and could fail to resend routes that the neighbor sent before rebooting.
- **CAM optimization for the default route and trunk groups** – The aggregation feature described in “Optimizing the Forwarding Cache for the Default Route” on page 42 did not work properly with trunk groups.
- **CLI** – In configurations where the startup-config file contained a command defining an encrypted OSPF authentication key, the CLI did not correctly read the command from the file, and consequently did not execute the command.
- **CLI** – In configurations where rate limiting was applied to a port and an ACL was used as part of the rate limiting configuration, the CLI did not accept the rate limiting command when reading the command from a startup-config file.
- **CLI** – The software did not display a message on the console to indicate when a module’s temperature had become higher than the configured warning level. In previous software releases and this release, the software does display a message if the shutdown temperature is reached. In addition, this release and previous releases generate Syslog messages and SNMP traps to indicate when a module has reached its warning or shutdown temperature.
- **CLI** – In some cases, the software did not update a virtual routing interface (VE) to reflect a change in the port membership of the VLAN containing the virtual interface. This could occur after a Telnet management session was established with the device.
- **CLI** – The software did not distinguish between normal collisions and late collisions on 10 Mbps ports. Software release 07.1.22 makes this distinction by adding a new field, OutLateCollisions. See “New Packet Error Statistic” on page 33.
- **CLI** – The software counted runt packets detected on 10 Mbps ports running in half-duplex mode as errors in

the InErrors field of the detailed **show statistics** display for individual ports. In software release 07.1.22, the software counts the collisions detected on this type of port in either the OutCollisions field or the OutLateCollisions field. See “New Packet Error Statistic” on page 33.

- **DHCP** – If a DHCP client was moved from one port-based VLAN to another, the client was unable to obtain an IP address.
- **DVMRP** and **PIM-DM** – In a configuration where two clients were in separate port-based VLANs that shared a tagged port, a client joining or leaving the multicast group could cause video interruption for the other client.
- **DVMRP** (chassis routing switch only) – In some configurations, if you ran DVMRP with multiple tagged trunk groups, the multicast packets were dropped when a new trunk member port was added to the trunk group.
- **DVMRP** – In a configuration where a client and server were in separate VLANs that shared the same tagged port, the device sometimes experienced very high CPU utilization or reloaded.
- **Hot swapping** (9304M/9308M only) – If 100 Mbps ports on a module were configured to operate at full-duplex, and the module was then swapped out, the ports on the module no longer operated in full-duplex mode.
- **IGMP** – If a client left a multicast group, other clients in the same port-based VLAN experienced video interruption.
- **IP multicast** – If the device was sending a lot of multicast traffic to a Layer 2 Switch, the routing switch’s software reset if the Layer 2 Switch was reset or the link between the devices was removed.
- **IP multicast** – If the device received a lot of multicast traffic while the software was loading, this caused the software to begin reloading again.
- **IP multicast** – On a routing switch running PIM, the prune state was not cleaned when a link went down. As a result, the routing switch could run out of prune state resources, causing a system reset.
- **NAT** – Heavy NAT traffic could cause the software to reload due to data corruption while the device was processing routes pending in the queue.
- **NAT** – Hardware buffer resources could become unavailable after several days of heavy NAT traffic.
- **OSPF** – This software did not delete LSAs associated with an OSPF interface that went down, which could result in routing loops.
- **OSPF** – The software advertised individual routes along with a summary route when static routes were summarized.
- **OSPF** – Disabling external route summarization did not result in the summary route being removed from a neighbor’s route table until the IP route table was cleared (**clear ip route** command).
- **OSPF** – When a route was advertised as an intra-area route as well as an external route from a non-backbone area, removing the intra-area route advertisement did not cause the corresponding type-3 LSA in the backbone area to be flushed from Area Border Routers (ABRs).
- **PIM DM** – In a configuration where a PIM DM client was attached to a tagged port that was a member of multiple port-based VLANs, when the client left a multicast group, traffic in the other VLANs sharing the tagged port could be interrupted until receipt of the next IGMP query.
- **Route-only option** – The software did not properly clear CAM resources when Layer 2 support was enabled or disabled on a port.
- **SNMP** – If you used SNMP to create an IP access policy for permitting or denying traffic, the software instead created a QoS policy.
- **Syslog** (9304M/9308M only) – The Syslog buffer could become full of “ARP mapping deleted“ messages although routing was functioning normally.
- **TCP counter** – If a TCP connection to the routing switch remained active for more than 49 days, this could cause a software reload.
- **Trunk groups** (9304M/9308M only) – If a MAC address was moved from one port to another into a trunk group, the device did not detect the change for one or two minutes. This caused Layer 2 traffic addressed to

the moved MAC address to be misforwarded until the device detected the change. This occurred regardless of whether the port the MAC address was moved from was also configured to be part of a trunk group. This issue did not occur if the MAC address was moved from one port to another within the same trunk group.

- **VRRP** – If a module on which a VRRP interface was configured was hot swapped (removed and then re-inserted), the VRRP address (the address being backed up) would not respond to IP pings.

Fixed in 07.1.19

This release applies to the HP 9304M and 9308M routing switches using an MII or Mini-GBIC (MIV) management module; that is, with redundant management.

- **Accessing the CLI** – If you attempted to access the CLI using RADIUS or TACACS+ security, or SSH, you were automatically placed at the Privileged EXEC level, rather than the User EXEC level.
- **Accessing the CLI** – If the option to suppress Telnet rejection messages was enabled (**telnet server suppress-reject-message** command), the device did not respond to TCP SYNs if the Telnet request was denied.
- **ACLs** (9304M and 9308M routing switches only) – A deny ACL could sometimes permit packets that the ACL was supposed to deny.
- **ACLs** – If you applied an ACL to an interface that had a QoS priority higher than 0 (the default), the device did not perform QoS for Layer 4 sessions on the interface. In addition, if you applied an ACL to a virtual interface and any of its physical interfaces had a QoS priority higher than 0, the device did not perform QoS for Layer 4 sessions on the interface.
- **ACLs and NAT** – The software did not translate a private IP address into a real IP address if an ACL was applied to an interface on which NAT was enabled.

NOTE: The software issue has been fixed, but make sure you also follow the guidelines described in “Using ACLs and Network Address Translation (NAT) on the Same Interface” on page 15.

- **BGP4** – If the routing switch was configured with route-flap dampening and inbound route maps that changed the attributes of BGP routes received from external neighbors, the software could reload after running for awhile.
- **BGP4** – If a BGP4 neighbor session on the routing switch was inactive and the routing switch was configured to use a password to authenticate sessions with the neighbor, the routing switch sometimes ended a Telnet management session after five seconds. This could occur when the neighbor session was inactive due to a password mismatch or when the BGP session was not configured on the neighbor.
- **CLI** – The **show ip bgp** command listed a route as the best route even when the route was unreachable. The show ip bgp route correctly displayed the route information.
- **CLI** – The CLI allowed you to specify an extended ACL ID (100 – 199) when configuring an ACL to secure management access though Telnet, even though this feature requires a standard ACL (1 – 99).
- **CLI** – The CLI could accept a duplicate Adaptive Rate Limiting policy even though the CLI displayed a warning message when the command was entered.
- **CLI** – The CLI did not allow you to apply an Adaptive Rate Limiting using a rate limiting ACL unless you first configured an IP ACL with the same ACL number.
- **CLI** – The **show ip interface** command listed the state of a port as up even though the port did not have a network cable plugged in.
- **CLI and Web management interface** – When a Gigabit interface was brought down and then brought back up, the software incorrectly showed 100% utilization for the port. This issue did not affect other types of ports.
- **IGMP** – The software did not time out group membership after a group’s client left the group. Instead, the software continued sending group traffic to the client.
- **Interface statistics** – If a Gigabit link went down, the octets counters for the port were reset to zero.

- **IPX** – The software would not forward IPX requests out of physical ports. This problem did not affect virtual interfaces.
- **MAC age timer** – If you set the MAC age timer to 80 seconds, the device actually aged the MAC table entries out after only 10 – 20 seconds.
- **NAT** – The feature did not work if you enabled inside NAT on a virtual interface.
- **NAT** – When the routing switch was under heavy load, the device reloaded if the software was performing NAT while there were routes pending in the queue.
- **NAT** – The feature did not translate private addresses into Internet addresses if an ACL for outbound traffic was applied to any interface on the device.
- **OSPF** – The software did not encrypt OSPF authentication strings by default when you saved the configuration. Instead, when you saved the configuration, the software saved the OSPF authentication strings as clear text by default.
- **OSPF** – If OSPF was enabled but no areas were configured, unplugging a cable from a port caused the software to reload.
- **OSPF** – An adjacency with an OSPF neighbor that was a Cisco Systems router could occasionally remain in the Exchange state indefinitely.
- **OSPF** – In situations where outbound Adaptive Rate Limiting was configured on an interface and the CPU utilization was above 90%, the software dropped OSPF neighbor sessions. This problem also could occur if the interface the OSPF neighbor session was on was configured with inbound Adaptive Rate Limiting.
- **PIM Sparse** – The software sent join messages to the Rendezvous Point (RP) even when the routing switch wasn't the Designated Router (DR).
- **PIM Sparse** – If multiple routing switches configured for PIM Sparse within the same LAN had the same preference value and also the same metric to a given source, the +es could not complete the assert process. A a result, an assert loser continued to forward multicast traffic.
- **SNMP** – The software did not support loading an ACL configuration file from a TFTP server into the running-config.
- **SNMP** – If single-instance STP was enabled on a device, the Bridge MIB returned values for single STP instead of returning STP values for the default VLAN. In addition, if the HP device had the root port, the dot1dStpRootPort and snVLanByPortStpRootPort objects returned the value 1 instead of 0, which they should have returned.
- **SSH** – Consecutive failed login attempts could cause the software to reload.
- **TACACS/TACACS+** – The software bypassed authentication of the enable password if you logged in with a superuser password and TACACS/TACACS+ authentication was used to authenticate the login attempt.
- **TACACS+** – Buffers for the TACACS+ subsystem filled to capacity, preventing further AAA operations.
- **Telnet CLI** – If you entered the **show configuration** command from a Telnet session to the CLI, the message "INFO: config data, primary copy checksum failed, try to read from backup" was displayed.
- **Virtual interfaces** – If you configured a virtual interface, then disabled it before configuring an IP address on the interface, the **enable** command did not re-enable the interface. Re-enabling the interface required removal and re-insertion of the physical link(s) for the interface.
- **VRRP** – In a configuration where VRRP was running on multiple ports on different forwarding modules, if one of the modules was removed and then reinserted, the device did not relearn the VRRP virtual MAC address. This caused the active device to stop forwarding packets.
- **Web management interface** – When you logged out of the Web management interface, you were prompted to save the configuration file, even if you had not made any changes to it.
- **Web management interface** – The interface did not support IPX SAP ACLs.
- **Web management interface** – Selecting the Monitor->IPX->Port Counter link could cause the software to reload.

- **Web management interface** – If a BGP4 route did not have an AS path, using the Web management interface to display the route could cause the software to reload.

Fixed in 07.1.10

- **10/100 ports** – When a 10/100 port was disabled, the link LED did not go dark and the port on the other end of the link did not indicate that the link was down.
- **ACL** – An extended ACL for IP protocol TCP or UDP did not take effect. The CLI allowed the ACL to be entered, but the ACL did not take effect and was not displayed in the running-config or in the **show ip acl** display.
- **ACLs** – Interfaces on which SRP was running and also ACL entries were configured blocked all outbound TCP packets regardless of the actual filter conditions of the ACL entries.
- **ACLs** – If you used an external configuration file to load ACLs and an **access-list** command in the file had a blank space in front of the command, the system reset when you loaded the configuration file. This occurred if you loaded the file from a TFTP server.
- **ACLs** – If you downloaded an ACL configuration file to the device's running-config and the file contained a **no access-list** <acl-id> command, the CLI displayed an “Error: No such entry” message and the device eventually reloaded.
- **ACLs** – When an ACL was applied to an interface, the data buffer containing a packet denied by the ACL could be freed twice.
- **ACLs** – If you configured an ACL with the IP address value 0.0.0.0 (equivalent to “any”), the software assumed the comparison mask also was 0.0.0.0 (“any”), regardless of the mask value you actually specified.
- **ACLs** – The standard ACL mode for TCP and UDP packets could result in some packets being forwarded that should have been denied. For example, if an ACL permitted source address 10.2.0.0 255.255.0.0 (ACL entry 10.2.0.0 0.0.255.255) but denied all other addresses, the software also allowed packets with source address 10.3.0.0 255.255.0.0. To prevent this behavior, you can use the strict ACL TCP mode and strict ACL UDP mode. These modes enable tighter control by sending each TCP and UDP packet to the CPU for ACL comparison. The strict ACL TCP mode and the strict ACL UDP mode are new in software release 07.1.10.
- **ACLs** – If you accidentally tried to delete a named ACL that did not exist, the device reset itself.
- **Adaptive Rate Limiting** – The CLI allowed the Excess Burst Size value only to be greater than the Normal Burst Size. This software release allows you to specify an Excess Burst Size that is equal to or greater than the Normal Burst Size.
- **Adaptive Rate Limiting** – The rate limiting commands for an interface appeared in the running-config before the IP address information for that interface. If you saved the configuration to the startup-config file, then reloaded the software, the Adaptive Rate Limiting feature displayed an error message stating that the IP address was not configured. As a result, the rate limiting did not take place.
- **Adaptive Rate Limiting** – The feature applied rate limiting to all IP traffic, including BGP4 and OSPF control packets and broadcast messages, which by default should not be rate limited.
- **Aging of Layer 4 session entries** – In some cases, a problem in the aging mechanism for Layer 4 session entries could cause the system to reset. This problem generally was associated with ACLs.
- **AppleTalk** – The software did not delete cached AppleTalk ARP entries following a Layer 2 topology change.
- **AppleTalk** – The routing switch did not respond to GetMyZone packets, which are used by some AppleTalk devices to obtain their home zones.
- **AppleTalk** – If you deleted a zone, the routing switch did not update zone information to reflect the deletion until you reloaded the software.
- **BGP** – The software did not properly interpret regular expressions that contained both an underscore (_), which matches on the end of an input string and other items, and a dollar sign ($), which also matches on the end of an input string. For example, the expression “(27_)+$" should match on “27”, “27 27", “27 27 27", and so on. However, the underscore had already matched on the end of the input string, so there was nothing left for the dollar sign to match on.

- **BGP4** – If the router received a BGP4 route with a very large AS path, the BGP4 connection with the neighbor that sent the route changed from the ESTABLISHED state to the INITIALIZE state.
- **BGP4** – In a configuration where AS-path filters were in use and the routing switch received a route containing a very large number of AS numbers (50) in one path attribute, the software could reset.
- **BGP4** – The first time BGP4 was enabled on a device, the BGP4 timer was not properly initialized. This required you to save the configuration and reload the software to initialize the timer. In software release 07.1.10, you do not need to reload the software to initialize the timer. The timer is properly initialized as soon as you enable BGP4.
- **BGP4** – The BGP4 Multi-Exit Discriminator (MED) attribute was not correctly handled in some situations. Releases earlier than 07.1.10 always passed the MED attribute to BGP4 neighbors, including EBGP neighbors. This was true even though the software did provide a workaround using the set metric command in route maps. The BGP4 MED attribute is now handled as follows:
  - If the MED is received from any neighbor (EBGP, IBGP, or confederation EBGP), the software can pass the MED to other IBGP and confederation EBGP neighbors.
  - If the MED is received from an EBGP neighbor, the software cannot pass the MED to other EBGP neighbors.
  - If the MED is received from an IBGP or confederation EBGP neighbor, the software can pass the MED to other EBGP neighbors so long as the BGP route is originated locally in the AS or confederation. The software determines this by checking to see whether the external AS-Path length is zero.

  In addition, you still can use the **set metric** command in a route map to change the MED for routes sent to or received from EBGP neighbors.
- **BGP4** – For BGP4 routes that were locally originated using the **network** command, the software sometimes might not update the NEXT-HOP and MED values when the corresponding IGP route in the IP route table was changed.
- **CLI** – If you entered a very long string when prompted for a Telnet password, then pressed Enter before the software timed out the access attempt, the device reset.
- **CLI** – The **show ip bgp neighbor** <ip-addr> **advertised-route [detail]** command did not correctly display the actual NEXT_HOP and COMMUNITIES attributes of a route sent to the neighbor.
- **CLI** – The **show default values** command listed the default ARP age as 20 minutes but should have listed the value as ten minutes. In the current software release, the command lists the correct value.
- **CLI** – The **show ip** command listed the default value for IP Proxy ARP as disabled but the **show default** command listed the default state for this parameter as enabled. IP Proxy ARP is disabled by default. The **show default** display has been changed to reflect this.
- **CLI** – The CLI limited the number of VRRP VRIDs that could be displayed.
- **CLI** – When the skip-page mode was enabled, the last page of a **show vlan** display was missing a few lines of data. In addition, if the command was entered repeatedly, the CLI displayed the message “all 13 display buffers are busy, please try later” and did not display the VLAN data.
- **CLI** – The interface configuration level IPX commands **ipx-rip-update-hop-count-increment** and **ipx-sap-update-hip-count-increment** were documented but were not present in the CLI. The commands are now present at the Interface configuration level.
- **Forwarding performance** – Devices with many entries in the ARP cache or IP forwarding cache could experience slowed performance due to high CPU utilization.
- **Gigabit autonegotiation** – If an HP Gigabit port was connected to a Cabletron Gigabit port, and the Cabletron port's receive line was disconnected, the HP device continued to report the link to be up.
- **IGMP** – The software did not save the **ip igmp query-interval** or **ip igmp max-response-time** command in the startup-config file, and thus did not reinstate the commands following a software reload.

NOTE: Make sure IP multicast routing is enabled before you configure IGMP parameters on a routing switch.

- **IP** – Momentary high CPU utilization could occur if the device had active IP static routes and was waiting for an ARP response from the next-hop gateway used by the static routes.

NOTE: This problem did not affect links between 9304M, 9308M, 6308M-SX, or 6208M-SX devices.

- **IP** – If you removed a secondary IP address, an IP cache entry associated with the address could be removed only by reloading the software. Until the cache entry was removed by reloading the software, the device could not forward traffic if the destination's next hop was the removed secondary address.
- **IP** – Proxy ARP was enabled by default but should have been disabled by default.
- **IP directed broadcast** – If the device had an empty startup-config file, IP directed broadcast forwarding was enabled by default. Normally, the feature is disabled by default. This problem did not occur on devices whose startup-config files contained CLI commands.
- **IP access policies** – When an IP access-policy was configured on an interface, and the IP cache contained a large number of entries, the routing switch could be unresponsive for seconds upon learning a new route while flushing IP cache and flow entries.
- **IP forwarding** – Forwarding entries for a next-hop router were corrupted.
- **IP forwarding** – In some cases, packets could be forwarded using a less specific route instead of a more specific route. Normally, when the device has multiple routes to a destination, the device selects the more specific route.
- **IP forwarding** – A transit packet with an invalid Router Alert option (with length 0) caused the router to hang.
- **IP forwarding** – The device did not properly send traffic from its IP stack when using an IP default network. This problem did not affect IP transit traffic.
- **IP Multicast** – In configurations that use tagged VLAN ports, a multicast group was not visible on different VLAN sub-nets after the group received its routing information from a router.
- **IP static routes** – In some configurations, the software did not install a static route in the IP route table when a direct route to the same destination went away.
- **IPX** – In configurations where a stream of IPX traffic (in one interface and out another) was using the same encapsulation type, if an IPX interface with a different encapsulation type was added, the older traffic stream was disrupted, causing the connection to time out. For example, if a device was receiving IPX 802.2 traffic on one interface and forwarding it to another interface using the same encapsulation type, this traffic was disrupted if a third interface using another encapsulation type (example: 802.3) was added.
- **Management module switchover** – Ports could experience packet loss following failover from the active to standby module.
- **Management module switchover** – During switchover, the new active module did not load the running-config from the other module but instead loaded the system using the new active module’s own startup-config file.
- **NAT** – In a configuration where NAT was configured on a virtual routing interface (“virtual interface” or “VE”), packet addresses were not translated between the inside and outside addresses.
- **OSPF** – In a network where routing topology frequently changed, disabling an OSPF interface could occasionally cause the routing switch to reset.
- **OSPF** – Inter-area routes could be created with the wrong next-hop address. This could occur when there were multiple ABRs between the backbone area and a non-backbone area. When this occurred, inter-area summary LSAs sometimes were not prevented from being flooded back into the backbone area after being flooded into a non-backbone area.
- **OSPF** – On routing switches running both OSPF and IPX, IPX could cause OSPF packet corruption.

- **OSPF** – This problem affected only configurations where two ASBRs each advertised a static route (redistributed into OSPF on the ASBRs) to the same external network, and where the advertisements resulted in other OSPF routers having two equal-cost paths to the external network. If the static route referred to an interface on one of the ASBRs as its next hop, and that interface flapped (went down and then came back up), one of the equal-cost paths was missing in the routers that received the static route advertisement from the ASBRs.
- **OSPF** – In configurations where there was more than one route to a stub network and the routes were though different next-hop routers, the software did not always choose the route with the shorter path. When this occurred, it was usually when the route with the shorter path flapped (went down and came back up).
- **OSPF** – In some situations, routes were not added to the IP route table correctly following a topology change. This could occur in the following situations:
  - If a route was advertised by more than one router, and some routers advertised the route as an external LSA while other routers advertised the route as an intra-area LSA.
  - If a route was advertised by more than one router, and some routers advertised the route as an external LSA with the host-bits set (for example, 22.22.22.255 instead of 22.22.22.0), while other routers advertised the route as a normal external LSA.
- **OSPF** – OSPF protocol update packets whose size was equal to the MTU size were not transmitted properly. This issue did not affect packets of any other size.
- **OSPF** – In some instances, calculation of inter-area routes resulted in regeneration of a summary LSA for routes that actually had not changed since the last summary LSA.
- **OSPF** – If you used the option to set the metric on routes redistributed into OSPF, the software did not apply the new metric. This issue affected static, directly attached, and BGP4 routes redistributed into OSPF.
- **OSPF** – If the routing switch found a lower cost intra-area path for a destination network and the routing switch already had a higher cost intra-area path for the same destination network, the routing switch added the lower cost path as an additional path and treated the two paths as equal cost load sharing paths (with the lower cost). In the current software release, the routing switch adds the lower cost path and discards the higher cost path.
- **OSPF** – In a configuration where there was more than one route to a stub network, if the best route (the route with the lowest cost) became unavailable, the software did not use another, available route to the stub network.
- **OSPF** – In configurations with two or more equal-cost Area Border Routers (ABRs), the routing switch could fail to remove the corresponding route path for external routes or inter-area summary routes when the link to one of these ABRs went down.
- **OSPF** – Removing a static default route could cause the device to reset. This occurred in configurations where the HP routing switch was configured with default-information originate set to always and the routing switch and other OSPF routers had achieved full adjacency.
- **OSPF** – The device did not properly relearn a Type-5 External OSPF route with a non-null forwarding address unless you manually cleared the route from the IP route table.
- **PIM** – In configurations where PIM was running on a VLAN interface, if a client sent a leave message to leave a multicast group, the software did not process the request properly. As a result, the client continued to receive multicast data from the group until the group expired.
- **PIM Dense** – On a VLAN containing tagged ports, the group reports received on a tagged port were not processed correctly. As a result, the tagged ports could be omitted from the forwarding entry, which could result in incorrect forwarding of multicast traffic.
- **PIM Sparse** – In some network topologies, a routing switch running PIM Sparse for a long time with heavy traffic stopped forwarding.
- **PIM Sparse** – A memory management issue could cause the routing switch to drop IP multicast packets.
- **PIM Sparse** – The timer entries were not scheduled correctly, which could result in timer-related PIM Sparse events and messages to occur at times other than when expected.

- **RADIUS** – RADIUS authentication stopped working after 256 authentications. As part of standard RADIUS operation, the RADIUS 8-bit sequence number rolls back to 0 after 255. However, the HP device was using a 16-bit counter for the authentications and thus expected 256 (0x1ff), whereas the sequence number received from the RADIUS server was 0 (which was correct).
- **RADIUS** – If multiple users tried to log in to the HP device at the same time, this could cause the HP device to be unable to send RADIUS request packets to the RADIUS server. When this occurred, the problem persisted until the HP device was rebooted.
- **Route maps and redistribution** – Redistribution of routes into BGP4 using route maps to filter route tag values did not work porperly.
- **Route-only option** – The route-only option was not enabled by default.
- **SNMP** – An SNMP get-Next for the MIB object snBgp4RouteOperStatusAsPathList caused the device to reload if the object did not have an AS path.
- **SNMP** – The MIB object dot1dTpFdbTable contained MAC entries for all VLANs.
- **SNMP and Syslog** – If you disabled a trap, the equivalent Syslog message remained enabled. For example, if you disabled an OSPF trap, the software still sent Syslog messages if the event that initiated the trap occurred, although the software did not send a trap.
- **Software upgrade** – If you upgraded a device running AppleTalk on an 06.6.x version of software to 07.1.x, then entered the **span** command for STP, the device reset. In addition, single-instance STP commands disappeared from the configuration after the software was upgraded.
- **Spanning Tree** – If you enabled single STP, saved the configuration, then reloaded, STP was disabled on the device following the reload.
- **SRP** – In configurations where IP clients used SNAP encapsulation instead of Ethernet II encapsulation, the clients could ping the real IP address of the backed up gateway but could not ping the virtual IP address of the backed up gateway.
- **SRP** – On random occasions, when the active routing switch (primary) was powered down, then powered back up, network connectivity was lost to the hosts connected to the primary routing switch for approximately one minute.
- **SSH** – SSH did not respond to authentication agent forwarding request, as described in the SSH Connection Protocol specification (section 4.5.1). This prevented certain clients that use this option from establishing SSH sessions with an HP device.
- **STP** – If STP was enabled on a protocol VLAN that included a trunk group, ports forwarded traffic even thought they should have been blocked by STP.
- **STP** – Global STP did not allow the priority to be changed unless a VLAN was configured on the device.
- **STP** – If single-instance STP was enabled, some ports cycled through the listening, learning, and blocking states even though they should have been in the blocking state.
- **STP** (single-instance STP only) – In configurations where single-instance STP was enabled on the routing switch, the device was the root bridge, and the ports connecting the device to other devices were tagged ports, the device sometimes sent incorrect information in the STP BPDUs or sent them out the wrong ports. This could cause the STP state on the ports of the other devices to sometimes be incorrect (Blocking when they should have been Forwarding, or Forwarding when they should have been Blocking).
- **STP** (single-instance STP only) – If you enabled single-instance STP on a device, then disabled the feature, the device dropped all BPDUs. As a result, the device did not set the STP state of its ports correctly.
- **Syslog** – The **no snmp-server enable traps ospf** command would only disable OSPF traps, not OSPF Syslog messages.

- **TACACS+ Authorization** – A user being authenticated by a TACACS+ server may have received read-only access, even though the user was authorized to receive port-config or super-user access. This occurred because the login process, comprised of authentication (to match a user's name and password) and exec authorization (to determine a user's privilege level), used the same source port value for both authentication and exec authorization. If the TACACS+ server sent a TCP RST packet during authentication, it may have had an effect on exec authorization, causing the user to receive the default privilege level (read-only).
- **TCP** – A device reloaded after running continuously for several months due to a seldom-occurring software problem related to TCP connections that were closed due to unusual network conditions.
- **Telnet** – In a configuration where two routing switches were directly attached by the same physical link but each side of the link was on a separate network, and each of the routing switches was configured with a static route that pointed to the other routing switch, the devices could not establish Telnet connections with one another even though they could respond to IP pings from one another.
- **TFTP** – If you tried to load more than 32 ACLs into the running-config by using TFTP to transfer the ACLs, the device reloaded.
- **Traceroute** – The HP device did not respond to route traces from some third-party devices to a loopback interface with sub-net mask 255.255.255.255 on an HP device.
- **Trunk groups** – In configurations where SRP, trunk groups, and Spanning Tree all were configured, ports did not properly learn the MAC address for the new root bridge following a topology change. As a result, loops could occur in the network.
- **Trunk groups** – In a configuration that contained trunk groups and a Layer 3 protocol VLAN, connectivity on the routing switch's first port could be disrupted due to a software problem.
- **Trunk groups** – In a trunk group containing more than two ports, if the links in the trunk went down, the device could sometimes switch the trunk traffic to the incorrect ports.
- **Trunk ports** – If all the links in a trunk group on an HP device went down or the device was reloaded, the forwarding information on the HP device at the other end of the trunk links was not correctly updated. As a result, the device was not able to properly forward traffic that the device would normally send over the unavailable trunk links.
- **Trunk ports** – On a pair of trunk ports configured as an IPX interface, if the secondary port in the trunk group became unavailable while it was forwarding traffic for the IPX interface, the traffic did not fail over to the primary port.
- **Unknown unicast limiting** – If a port was configured for unknown unicast limiting, the limit was applied to all ports on the device.
- **VRRP** – During heavy traffic loads, the Backup VRRP router sometimes prematurely transitioned to Master, then returned to Backup soon after that due to failing to receive VRRP messages within the dead interval.
- **VRRP** – VRID priority settings were not displayed in the running-config.
- **VRRP** – If you deleted an IP address from an interface on which multiple VRIDs were configured, the software removed all the VRIDs in addition to the one that matched the deleted IP address.
- **VRRP** – If a device running VRRP in the backup state received a packet with the destination MAC of the VRID, the device tried to route the packet instead of forwarding it at Layer 2 to the VRRP master.
- **Web management interface** – The Port display panel showed the wrong port numbers for trunk ports. For example, if you configured ports 13 and 14 as trunk ports, the Port panel showed ports 12 and 13 as green (active) trunk ports. This was a display issue only and did not affect the operation of the trunk ports.
- **Web management interface** – If you were using the NetScape browser and enabled the front panel display, the browser would hang and not download all the required files.
- **Web management interface** – The interface did not allow creation of an AppleTalk protocol VLAN. The appropriate radio button could be selected, but selecting Add after selecting the radio button resulted in an error message.

Known Issues

Known Issues in Release 07.1.24

- **Mini-GBIC ports** – Hewlett-Packard offers and supports only mini-GBICs that include an HP label (with product number J4858A or J4859A) for use with the J4856A HP Procurve 9300 Mini-GBIC Module and the J4857A HP Procurve 9300 Mini-GBIC Redundant Management Module. Use of other brands of mini-GBICs is not supported.
- **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.
- **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.
- **SNMP** — The SNMP agent cannot support virtual interface numbers higher than 255.

Single STP Issues When Migrating from 06.6.x to 07.1.x

One of the enhancements in software release 07.1.x is support for up to 4095 port-based VLANs in a single spanning tree. This support required a change to the position of the port-based VLAN commands in the running-config and startup-config file.

- In software release 06.6.x, the VLAN commands are placed before the **spanning-tree single** command.
- In software release 07.1.x, the **spanning-tree single** command is placed before the VLAN configuration commands.

As a result of the changed command positions, if you boot a device using software release 07.1.x but also load a startup-config file created using software release 06.6.x, the CLI parser does not find the **spanning-tree single** command before the VLAN commands. The parser therefore assumes that the single STP feature is not enabled. When the device finishes booting, the device contains a separate spanning tree for each VLAN on which STP is enabled, instead of a single spanning tree consisting of all the VLANs on which STP is enabled.

Migration Procedure

NOTE: You need to use this procedure only if you are upgrading a device running software release 06.6.x and using single STP to software release 07.1.x.

To migrate your single STP configuration from 06.6.x to 07.7.x:

- Make a backup copy of the startup-config file. You will need this file if you decide to revert to the 06.6.x release for any reason.
- Boot the device using software release 07.1.x.
- Disable single STP.
- Enable STP (not single STP) in each of the port-based VLANs that you want to include in the single spanning tree.
- Enable single STP.
- Save the configuration. You cannot use the configuration you saved using 06.6.x on a device running 07.1.x.

Saving a Backup Copy of the Service's Startup-Config File

1. Make sure the device has IP access to a TFTP server.
2. Enter one of the following commands at the Privileged EXEC level of the CLI to copy the device's startup-config file onto the TFTP server:
   - **copy startup-config tftp** <ip-addr> <filename>
   - **ncopy startup-config tftp** <ip-addr> <from-name>

Completing the Migration

1. Boot the device using software release 07.1.x.
2. Enter the following command at the global CONFIG level of the CLI to disable single STP:
   - **no spanning-tree single**
3. Enable STP within the port-based VLANs that will be members of the single spanning tree. When you re-enable single STP, all the VLANs in which you enabled STP will become members of the single spanning tree. Other VLANs (in which STP is disabled), will not become part of the single spanning tree.

   To enable STP in a VLAN, enter the following command at the global CONFIG level of the CLI to exchange the CLI to the configuration level for that VLAN:
   - **vlan** <vlan-id>

   To enable STP within the VLAN, enter the following command:
   - **spanning-tree**
4. Enable single STP. To do so, enter the **exit** command to return to the global CONFIG level of the CLI, then enter the following command to enable single STP:
   - **spanning-tree single**
5. Save the configuration changes to the startup-config file.
6. Reload the 07.1.x software.

NOTE: When you reload, use the startup-config file you saved in Step 5. If you try to use a startup-config file saved while running 06.6.x, the single STP configuration will not be loaded.

HP Part Number: 5969-2390
Edition 1, September 2001