



Release Notes *for Version 06.6.28 and Version 07.1.10 Operating Systems*

for the HP ProCurve Routing Switch 9304M, 9308M, and 6308M-SX, and Switch 6208M-SX

These release notes describe:

- The new operating system enhancements not available in software releases prior to version 07.1.10 for the HP ProCurve 9304M and 9308M routing switches with Redundant Management (MII)
- Earlier software operating problems fixed in version 07.1.10 for the HP ProCurve 9304M and 9308M routing switches with Redundant Management (MII)
- Earlier software operating problems fixed in version 06.6.28 for the HP Procurve 9304M and 9308M routing switches without Redundant Management (MI), the HP ProCurve 6308M-SX routing switch, and the HP 6208M-SX switch

For release notes describing earlier software releases, go to the **technical support | manuals** area of the HP ProCurve website at <http://www.hp.com/go/hpprocurve>.

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

Contents

SOFTWARE BRANCHES	2
REDUNDANT MANAGEMENT ON THE 9304M AND 9308M ROUTING SWITCHES	2
DOWNLOADING SOFTWARE AND DOCUMENTATION	3
TO DOWNLOAD A SOFTWARE VERSION:	3
TO DOWNLOAD PRODUCT DOCUMENTATION:	3
SOFTWARE/DEVICE COMPATIBILITY	3
ALREADY USING A 9304M OR 9308M WITH REDUNDANT MANAGEMENT? HERE'S NEW INFORMATION!	4
SOFTWARE IMAGE FILES	5
NOTE REGARDING DISABLING BGP4, OSPF, OR VRRP	5
SUMMARY OF ENHANCEMENTS FOR SOFTWARE RELEASE 07.1.10 FOR DEVICES WITH	
REDUNDANT MANAGEMENT (MII) MODULES	6
SYSTEM LEVEL ENHANCEMENTS	6
LAYER 3 ENHANCEMENTS	6
LAYER 2 ENHANCEMENTS	8
SYSTEM LEVEL ENHANCEMENTS	9
SOFTWARE FIXES	11
FIXED IN 06.6.28	11
FIXED IN 07.1.10	13
KNOWN ISSUES IN RELEASES 06.6.28 AND 07.1.10	18

Software Branches

Beginning with the software releases covered in this document, 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 later 07.X and higher releases:** These releases typically may include new features, enhancements to existing features, and bug fixes, and operate only with the HP 9304M and 9308M routing switches with redundant management (MII modules).

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" modules (for Management 2"). MII modules include:

- J4845A ProCurve 9300 GigLX Redundant Management Module (8-port)
- J4846A ProCurve 9300 GigSX Redundant Management Module (8-port)
- J4847A ProCurve 9300 Redundant Management Module (0-port)

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

NOTE: MI and MII management modules are mutually exclusive. That is, the routing switch does not operate if both an MII management module and an MI management module are installed at the same time.

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>.

NOTE: The flash image files for this software release differ depending on the product. See "Software Image Files" on page 1.

NOTE: 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.

These notes also contain information regarding what happens when you disable BGP4, OSPF, or VRRP. See “Note Regarding Disabling BGP4, OSPF, or VRRP” on page 5.

Downloading Software and Documentation

Software versions 06.6.28 or 07.1.10 and the corresponding product documentation can be downloaded 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 new manuals under the heading “**For software version 06.6.28 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	HPR05216.BIN HPR06605.BIN HPR06616.BIN HPR06628.BIN	HPS05216.BIN HPS06605.BIN HPS06616.BIN HPS06628.BIN
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) <i>With Redundant Management Module(s) (MII)</i>	No	No	Yes	No	No
HP ProCurve Routing Switch 9304M (J4139A) and 9308M (J4138A) <i>Without Redundant Management Modules (MI)</i>	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

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

If you have a 9304M or 9308M routing switch that was shipped before 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 one of the above devices before software release 07.6.10 began shipping, and you are updating the device to release 07.1.10, then you may want to examine the new product manuals that are available beginning with the 07.1.10 release. To view (and freely download) PDF versions of these manuals (chapter-by-chapter files). See "To Download Product Documentation:" on page 3.

Software Image Files

To run either software release 06.6.28 or 07.1.10, you need the indicated boot and flash images listed in the following table.

Product	Boot Image	Flash Image
HP 9304M HP 9308M With one of these MI modules; that is, without Redundant Management: <ul style="list-style-type: none"> • J4140A • J4144A • J4146A) 	M1B07108.bin or later recommended	HPR06628.bin*
HP 9304M HP 9308M With any one or two of these MII modules; that is, with Redundant Management: <ul style="list-style-type: none"> • J4846A • J4845A • J4847A 	M2B07108.bin	H2R07110.bin
HP 6308M-SX	• M1B07108.bin or later recommended	• HPR06628.bin*
HP 6208M-SX	• M1B07108.bin or later recommended	• HPS06628.bin*

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

NOTE: If you are adding a Gigabit Copper module to a 9304M or 9308M chassis, you must upgrade to boot code version M1B07108.bin or later for MI devices, or M2B07108.bin or later for MII devices.

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.

Summary of Enhancements for Software Release 07.1.10 for Devices with Redundant Management (MII) Modules

This section summarizes the operating system enhancements in software release 07.1.10 for the HP 9304M and 9308M with Redundant Management (MII) modules installed. For details about the enhancements, refer to the documentation provided for release 06.6.28 and 07.1.10. If your MII routing switch management module was shipped with release 07.1.10, 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 4 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 13.

NOTE: These notes do not include descriptions of enhancements added in software releases after 06.6.16. For details about enhancements in earlier releases, see the release notes for the earlier releases.

Software release 07.1.10 contains the following enhancements.

System Level Enhancements

Software release 07.1.10 contains the following 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.

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. Note: This enhancement applies only to routing switches. The maximum number of static ARP entries supported depends on the amount of DRAM memory on the management module.
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 chassis-based routing switch to perform host-based IP load sharing for specific routes while performing network-based IP load sharing for the other routes.

Enhancement	Description
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.
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.

Enhancement	Description
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: <ul style="list-style-type: none"> • 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.

Layer 2 Enhancements

Software release 07.1.10 contains the following 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.

System Level Enhancements

Software release 07.1.10 contains the following system-level enhancements:

Enhancement	Description
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
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 modules. Virtual interface groups are supported only on the chassis-based 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.

Enhancement	Description
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.
POS Frame Relay	You can configure POS ports for Frame Relay.
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 file and the running-config: <ul style="list-style-type: none"> 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.
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.

Enhancement	Description
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.

Software Fixes

Fixed in 06.6.28

This section lists the problems that have been fixed in software release 06.6.28 for the following devices:

- HP 9304M and 9308M routing switches using an MI management module; that is, without redundant management
 - HP 6308M-SX routing switch
 - HP 6208M-SX switch
- . For information about fixes in a software release before 06.6.28, see the release notes for that release.
- **10/100 ports** (9304M and 9308M only) – 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.
 - **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 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 or a flash card.
 - **ACLs** – 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.
 - **ARP** – In configurations that use the IP follow feature, which allows multiple port-based VLANs to share the same IP sub-net address, ARP entries entered the Invalid state and remained in this state until the ARP entries were cleared. This problem prevented the device from responding to IP pings.
 - **BGP4** – If a route map without **set** commands for matched routes was used for filtering neighbor outbound routes, the route map could cause reference count errors for the BGP attributes. As a result, attributes learned from a neighbor might not be cleared from memory even when the attributes were no longer being used. This could cause the memory for attributes to become full.
 - **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 06.6.28, you do not need to reload the software to initialize the timer. The timer is properly initialized as soon as you enable BGP4.
 - **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.
 - **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** (6208M-SX switch only) – The interface erroneously stated that the software supported Secure Shell (SSH).

- **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** – If you entered the `ip pim ttl-threshold <num>` command, after you saved the configuration change to the startup-config file, the file contained two instances of the command. Moreover, after you reloaded the software, the `show ip pim int <portnum>` command showed the wrong TTL threshold value.
- **IGMP** (routing switch only) – 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.
- **MAC filters** – A MAC filter applied as Ethernet type 0800 and equal to 0806 did not work.
- **OSPF** – Removing a static default route caused the system to reset.
- **OSPF** – OSPF permit redistribute did not work properly.
- **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** – 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** – 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 through 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).
- **PIM and VLANs** – Outbound multicast packets on a tagged PIM interface were sent with the VLAN ID 0.
- **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** – 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).
- **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.
- **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.
- **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.
- **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.
- **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.
- **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.

Fixed in 07.1.10

This section lists the problems that have been fixed in software release 07.1.10 for HP 9304M and 9308M routing switches using an MII management module; that is, with redundant management. For information about fixes in an earlier software release, see the release notes for that release.

- **10/100 ports** (9304M and 9308M only) – 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** (9304M and 9308M routing switches only) – 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.

- **Aging of Layer 4 session entries** (9304M and 9308M routing switch only) – 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** (9304M and 9308M routing switch only) – 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. Software release 07.1.10 handles the BGP4 MED attribute 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.

- **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.

- **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 (6208M-SX switch only)** – A device running IGMP multicast containment (IP multicast traffic reduction) did not correctly discover multicast routers running PIM Dense mode when more than one such router was present within a VLAN. As a result, multicast streams could be interrupted.
- **IGMP (routing switch only)** – 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.

NOTE: This issue affected chassis-based routing switches only.

- **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 (9304M and 9308M only)** – 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 (9304M and 9308M routing switch only)** – 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.
- **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 through 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.

- **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.
- **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.
- **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 though 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 (chassis devices only)** – 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.
- **Syslog** – The `no snmp-server enable traps ospf` command would only disable OSPF traps, not OSPF Syslog messages.
- **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.
- **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** – The enhancement in 07.1.10 that automatically re-balanced traffic across trunk group links when a down link came back up did not work properly.
- **Trunk groups** – In a configuration that contained trunk groups and a Layer 3 protocol VLAN, connectivity on the Layer 3 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 (9304M and 9308M only)** – 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 in Releases 06.6.28 and 07.1.10

This software release contains the following issues.

- **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.
- **DVMRP (routing switches only)** – In some configurations, if you run DVMRP with multiple tagged trunk groups, the multicast packets are dropped when a new trunk member port is added to the trunk group.
- **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, `HP9300 RouterHP9300 Router>`). This is a cosmetic issue only and does not affect the device's operation or performance. This issue is also present on the 6208M-SX and 6308M-SX devices.
- **CLI** – The **show cpu** command sometimes displays incorrect statistics.
- **CLI (stackable devices only)** – If a Gigabit port is configured for Auto-Gigabit (auto-negotiation) in software release 06.6.16 or later, the device does not retain the change following a software reload. The running-config file contains a Negotiation-Off setting for the port and the port at the other end of the link does not come up.
- **Interface-Based Static Routes** – If you replace an interface-based static route with one that has a longer

network mask, the interface-based static route with the longer network mask does not correctly supersede the existing interface-based static route with the shorter network mask.



© 2001 Hewlett-Packard Company. All rights reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

HP Part Number: 5969-2370
Edition 1, January 2001

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

