Functional Restrictions for the Avaya P550R and P880 Multiservice Switches, *Software Version 5.3.2*

Overview

This document lists functional restrictions of Avaya P550R® and P880 Multiservice switches that are running version 5.3.2 application software. Version 5.3.2 of the application software is supported only by the M5501R-SUP (P550R) and M5500R-SUPA (P880) supervisor modules.

Functional restrictions are issues that restrict the functionality of a particular feature. For most of these issues a workaround exists. This document is periodically updated with more current information.

For detailed information about your product, see the following documents:

- *Installation Guide for the Avaya P550R, P580, P880, and P882 Multiservice Switches, Version 5.3.1*
- *Avaya P550R, P580, P880, and P882 Multiservice Switch User Guide, Version 5.3.1*
- *Command Reference Guide for the Avaya P550R, P580, P880, and P882 Multiservice Switches, Version 5.3.1*
- *Avaya Multiservice Switch Error Messages*
- *Release Notes for the Avaya P550R and P880 Multiservice Switches, Software Version 5.3.2*
- *User Guide Addendum for the Avaya P550R and P880 Multiservice Switches, Software Version 5.3.2*

To download software, the latest release notes, and other documentation, see http://support.avaya.com. Click **Product Documentation**.

*** Important:** Avaya strongly recommends that you run the latest version of application software on your switch.

This document includes the following topics:

Functional Restrictions and Workarounds

The following functional restrictions and workarounds apply to Avaya P550R and P880 Multiservice switches.

24-Port, 100BASE-FX Module

Issue **(RN000151)** The Physical Port Configuration Web page and **show port physical** CLI command inaccurately display a connector of Fiber SC for the 80-series, 24-port, 100BASE-FX module. The correct connector type is MT-RJ.

Workaround Not applicable.

50-Series, 48-Port, 10/100 Base-TX Module

Issue **(RN000241)** The Physical Port Configuration Web page for 50-Series, 48-port 10/100 Base-TX modules with Telco connectors (model number M5548E-100TC) displays **Unknown** in the **Type** column.

Workaround Use the CLI or SNMP MIB files to view the correct port type.

Address Forwarding Table

Issue **(RN000251)** The following error message is incomplete when it is displayed by the CLI and event log:

```
Undefined Error--FILE: pleaft.cpp LINE: 1595
Value 0x0: Hash Tbl at max size. New MAC hash
values picked, will be active upo
```

The last sentence should read:

```
New MAC hash values picked, will be active upon
reboot.
```

This error message is displayed only if Debug mode is enabled.

Workaround Not applicable.

Issue **(RN000152)** If you create a static entry in the AFT and associate the entry with a port on the 50-series, 48-port, 10/100Base-TX module, the **show aft entry** command displays the wrong port for the entry.

Workaround Use the Address Entry Search Web page to search the AFT. To access this Web page, expand the **L2 Switching** > **Address Forwarding Table** folders, and click **Address Search**.

AppleTalk

Issue **(RN000274)** If more than 20 AppleTalk interfaces are set up on the switch, you cannot delete an AppleTalk interface by using the Web Agent.

Workaround Use the **no interface <*intf-name*>** CLI command to delete AppleTalk interfaces.

Issue **(RN000247)** If you use the **ping appletalk** command to ping a local AppleTalk interface, the ping times out. However, remote AppleTalk interfaces reply to the command.

Workaround Use the **show appletalk interface** command to determine the status of local AppleTalk interfaces.

Issue **(RN000153)** When multiple zones are set up on an AppleTalk interface, the **no appletalk zone <*zone-name*>** CLI command does not delete a single zone.

Workaround Use the Web Agent to delete a single AppleTalk zone. Alternately, you can use the **no appletalk zone** command to delete all zones on the interface and then reassign zones to the interface.

ATM Uplink Module

Issue **(RN000220)** If a P880 has its supervisor module in slot 2 and the switch resets, the switch ports on the ATM Uplink modules do not reestablish connectivity. This issue occurs only if slot 1 is empty.

Workaround Reset the ATM Uplink module.

If the switch has a single supervisor module, Avaya recommends that you install it in slot 1. Doing so prevents this issue from occurring. If you need to move a supervisor module, move it during a maintenance window because you must turn off the switch.

Avaya MultiService Network Manager

Issue **(RN000262)** If you use the Avaya MultiService Network Manager, (MSNM), (formerly known as CajunView), and you create a user name on the switch that has the same name as the SNMP community string for that switch, when you attempt to log in using that user name, access will be denied.

Workaround Do not create a user name with the same name as the SNMP community string when you use MSNM.

Custom Access Types

Issue **(RN000272)** The Web Agent and CLI have different default settings for custom access types. When you create a new custom access type:

- The Web Agent, by default, assigns read-write permission for all features.
- The CLI, by default, assigns no access to any features.

Workaround Not applicable.

DNS

Issue **(RN000254)** The switch does not resolve partial domain name suffixes. For example, if **documentation.support.avaya.com**, **support.avaya.com**, and **avaya.com** are set as domain name suffixes on the switch, and you enter **ping documentation.support**, the switch does not resolve the domain name.

Workaround Always enter the fully qualified domain name.

Issue **(RN000003)** If you modify or delete a domain name suffix, the changes do not take effect until you reset the switch.

Workaround After modifying or deleting a domain name suffix, save the running configuration to the startup configuration and reset the switch during a scheduled maintenance period.

DVMRP

Issue **(RN000248)** If you change the default setting of the **Prune Message Retransmit Interval** field on the DVMRP Global Configuration Web page, the change does not take effect. The Avaya Multiservice switches do not currently support retransmission of prune messages.

Workaround Not applicable.

Issue **(RN000229)** If the switch is forwarding a multicast flow to multiple destinations in the network, two of which are reached by DVMRP tunnel, the switch may forward incorrect multicast packets to the destinations that are *not* reached by DVMRP tunnel.

For example if the switch is forwarding a multicast flow to destinations A, B, and C, and destinations A and B are being reached by DVMRP tunnel, destination C may receive incorrect multicast packets.

Workaround Some possible workarounds:

- Use only one DVMRP tunnel per switch.
- Use DVMRP tunnels for other destinations also.

This problem exist only when active multicast sessions simultaneously exist for broadcast interfaces and two DVMRP tunnels.

Issue **(RN000130)** The **mtrace** CLI command does not work. If you enter the command, no information is displayed.

Workaround You can obtain mtrace information by conducting an mtrace from another device in the multicast path. If you use this workaround, make sure that mtrace support is enabled on all routers in the path of the multicast flow.

You can enable mtrace support on Avaya Multiservice switches by using either the IGMP Global Configuration Web page or the **ip mtrace** CLI command.

Issue

(RN000131) The switch advertises DVMRP routes that have a state of **DOWN**. This issue is caused by the routing table falsely reporting **DOWN** interfaces as **UP** interfaces.

This issue occurs only when you change one of the DVMRP Global Configuration settings, including enabling DVMRP itself.

Workaround

Set **Multicast Routing Protocol** to **None** for interfaces that have a state of **DOWN**.

Issue

(RN000091) If you use the IP Interfaces Web page to simultaneously change the IP address of an interface and enable DVMRP on the interface, DVMP will not work correctly, and the interface will not support multicast routing.

Symptoms of this issue include:

- The interface does not accept IGMP membership reports.
- The CLI does not display the interface when you enter the **show ip igmp interface command**.
- The IGMP Interfaces Web page does not display the interface.

Workaround

To change the IP address of an interface and enable DVMRP on the interface:

1. Change the IP address of the interface.
2. Click **Apply**.
3. Enable DVMRP for the interface.
4. Click **Apply**.

Flood Rate Limiting

Issue

(RN000004) The options for the **Rate Limit Rate** parameter do not work correctly on:

- 80-series, 10/100 Ethernet modules
- 80-series, 100 Base-FX Ethernet modules

These options do not work correctly regardless of whether you set them by using the Web Agent or CLI. The default setting of 20% provides no rate limiting during broadcast storms.

Workaround

Use the following table to determine the option that corresponds to the rate that you want to set. Avaya recommends that you enter a rate of 2%, which corresponds to the default setting of 20%.

Table 1. Flood Rate Limiting on 80-Series 10/100 Modules

Desired Rate	Web or CLI Setting
80%	10%
40%	5%
20%	2%
10%	1%

80-Series Gigabit Modules

Issue

(RN000154) The switch does not display the number of transmit errors that occur on 80-series gigabit modules. The following Web pages and CLI commands display **0** for the number of transmit errors that have occurred, regardless of whether or not errors have occurred:

- Port Statistics for Module *x* Web page, **Errors** column for transmit traffic (**Modules & Ports** > **Statistics** > **Module *x***).
- Ethernet Interface Statistics for Port *x.y* Web page, **Errors** field in **Transmit** column of the Ethernet Interface Statistics table (**Modules & Ports** > **Statistics** > **Module *x*** > **Port *x.y***).
- **show port counters** CLI command, **Transmit Errors** field.
- **show ethernet counters** CLI command, **Errors** field in **Transmit** column and **TxCRC Errors** field.

Workaround

Not applicable.

Hunt Groups

Issue

(RN000270) The following issue occurs only on 50-series, 12-port 10/100Base-TX modules (M5512R-100TX) and 50-series, 20-port 10/100Base-TX modules (M5520-100TX).

If Spanning Tree is enabled on the switch, the switch may stop forwarding traffic to hunt group ports for 30 seconds or less if you:

- Disable or disconnect the flood port of a hunt group
- Remove a module that has a port participating in a hunt group (regardless of whether the port is the flood port).
- Reinsert a module that has a port participating in a hunt group. When you reinsert the module, flooding may also occur for up to 30 seconds.

Workaround

Enable Fast Start on the hunt group ports. You can enable this feature on the Switch Port Configuration Web page.

Issue

(RN000260) When you assign a static MAC address to the base port of a hunt group, the hunt group link stops transmitting. When you assign a static MAC address to a nonbase port, the switch functions correctly.

Workaround

Do not assign a static MAC address to the base port of a hunt group.

Issue

(RN000094) If you remove the module that has the base port for a hunt group, a loss of traffic will occur for approximately eight seconds. This problem occurs only if:

- The base port is on a licensed, 80-series module that is in an 80-series chassis

  AND

- All other ports on the switch that are participating in the hunt group are on licensed, 80-series modules.

During this 8-second period, VRRP mastership may change to the backup router and then back to the original active router. This changing of VRRP mastership may cause OSPF adjacencies and routes to be lost.

Workaround

Set the VRRP advertisement timer to 4 seconds or greater, and set OSPF timers so that the Dead interval is 12 seconds or greater. Alternatively, remove modules from chassis only during a maintenance period.

IGMP

Issue

(RN000156) If the switch is serving as a nonquerier for a specific interface, the switch does not remove entries from the IGMP Group Membership Table when they expire.

This issue has no adverse effect on the network unless the switch becomes the designated querier for the interface. If the switch were to become the designated querier, it would forward multicast flows for the entries in the IGMP Group Membership Table that should have been removed.

Workaround

Flush the IGMP Group Membership Table by clicking **Flush Table** on the IGMP Group Membership Table Web page.

Intelligent Multicasting

Issue

(RN000157) If you set up a piggyback port to mirror a router port for intelligent multicast, the switch deletes the router port, and then:

- If the switch learned the router port dynamically, the piggyback port is displayed as the router port when the switch relearns the router port.

  The Router Port Display/Configuration Web page and **show intelligent-multicast router-port** CLI command incorrectly display the piggyback port, not the router port, in the **Port** field.

- If you set up the router port manually, you must set up the router port again.

  You can use the Router Port Display/Configuration Web page or **set intelligent-multicast router-port vlan** CLI command to set up a static router port.

Workaround

Do not set up a piggyback port to mirror an intelligent multicast router port.

Issue

(RN000158) If you assign port 1 of a module to a hunt group and the module is in slot 9 of a P882 or P880, intelligent multicast traffic is not forwarded.

This issue occurs regardless of whether port 1 of the module is physically connected to the hunt group. Merely *assigning* the port to the hunt group causes the issue to occur.

Workaround

Do not assign port 1 of a module to a hunt group if the module is in slot 9 of a P882 or P880.

Issue

(RN000097) If an LGMP server and LGMP client are participating in a hunt group and the flood link fails, multicast traffic stops on the LGMP client for approximately 4 to 6 minutes. As soon as the current IGMP session expires and the LGMP server receives another IGMP membership report from the LGMP client, multicast traffic resumes.

Workaround

Not applicable.

Issue

(RN000016) Intelligent Multicasting can block protocols to nonmulticast routers. If you have enabled Intelligent Multicasting and have configured a VLAN to one or more nonmulticast routers or multicast-capable endstations, Intelligent Multicasting configures router ports where multicast-enabled routers reside. These multicast router ports allow all multicast packets to the adjacent multicast routers. Ports connected to nonmulticast-enabled routers are not considered router ports and do not receive multicast traffic. The issue can occur when multiple IP multicast addresses are mapped to the same multicast MAC address, resulting in protocol packets not being sent to the adjacent non-multicast enabled routers.

Example:

The unicast routing protocol in use on all connected routers is OSPF, and all ports are on the same VLAN. An endstation joins the IP multicast group 226.128.0.5 on port 1. The MAC address for the group is 01:00:5E:00:00:05. IGMP snooping creates a session for this MAC address, with port 1 as the client port. There is a non-multicast OSPF router attached to port 2. OSPF uses the IP multicast link scoped group 224.0.0.5, which also maps to a MAC address of 01:00:5E:00:00:05. Because port 2 is not considered a router port, and it is not part of the 01:00:5E:00:00:05 session, the switch only passes OSPF messages out port 1. Other protocols, such as the

Service Location Protocol (RFC 2608), use 224.0.1.22 and 224.0.1.35, which can be blocked by endstations joining sessions that map to the same MAC address.

Note: This is the same issue that is discussed in the Microsoft Product Support Article Q223136 involving RRAS setup. This specific issue, however, should not break routing protocols as suggested in this article because the Avaya multiservice switches ignore joins for local multicast groups (224.0.0.x).

Workaround

Check that all ports that are connected to the router are configured as router ports. This prevents all router-to-router messages from being blocked.

If other non-router protocols, such as the Server Location Service, are in use, create static sessions as needed. Do not create static sessions that conflict with the protocols used on your network.

Refer to the following Web site for a complete list of internet multicast addresses recognized by the IANA:

http://www.iana.org/assignments/multicast-addresses

*** Note:** **Enable** is the default setting for Rate Limiting on 10/100 Mbps ports. Multicast traffic and broadcast traffic is rate-limited (to 20%) on 10/100 Mbps ports.

Multicast traffic is rate-limited unless Intelligent Multicasting is enabled. Multicast traffic for which the Intelligent Multicast session was created is not subject to rate limiting unless the rate limiting state is set to **Enable** (all multicasts included).

Intrusion Traps

Issue

(RN000159) If intrusion traps are enabled on 80-series media modules, and you reset the switch, it stops sending intrusion traps for unknown source MAC addresses on those modules.

This issue occurs only on 80-series media modules.

Workaround

Disable and then reenable intrusion traps after you reset the switch. You can perform this task by using either the Switch Port Configuration Web page or **set port intrusion-trap** CLI command.

ip route Command

Issue

(RN000240) If you use the **ip route <*route-addr*> <*mask*> <*next-hop*> <*cost*>** command to create a default gateway, and a default gateway is already configured, the switch ignores the command and does not display an error message. If the default gateway is *not* present, the **ip route** command creates the default gateway. Avaya does not recommend that you use the **ip route** command to create or modify the default gateway.

Workaround

Use the command **ip default-gateway <*ip-address*>** to create or modify the default gateway.

IP Source Routing

Issue

(RN000161) If you disable IP source routing, the switch continues to allow source routing.

Workaround

Not applicable.

IPX

Issue

(RN000164) It is possible to create a maximum of only 22 IPX interfaces when you are using the Web Agent. If you have created 22 interfaces, it is not possible to delete any of the existing IPX interfaces by using the Web Agent.

Workaround

Use the CLI to configure or delete IPX interfaces when there are 22 or more interfaces.

Layer 3 Forwarding Cache

Issue

(RN000165) If you enter an invalid age interval on the Layer-3 Forwarding Cache Configuration Web page, the switch does not display an error message. When you reset the switch, the **Age interval** field reverts to the default setting, 120 seconds. Valid age intervals range from 20 to 360 seconds.

Workaround

Enter a valid age interval. Do not set the age interval to more than 360 seconds.

Multiservice Network Manager

Issue

(RN000236) When you use the P580/P882 Device Manager in Multiservice Network Manager (formerly called CajunView) to access ATM Uplink modules in the switch, incorrect data may be displayed. The incorrect values may be displayed when you use an SNMP request to access ATM module information through the ProminetMib.txt file.

Workaround

To obtain accurate module and port information for the ATM Uplink module, use the CLI or Web Agent.

OSPF

Issue

(RN000252) If the switch is a not-so-stubby-area (NSSA) router, it continuously generates a default NSSA link state advertisement (LSA). This continuous generation of an NSSA LSA causes the default route to be intermittently cleared from the routing table.

Because the NSSA router relies on the default route for all external traffic, this issue can cause a loss of traffic between the NSSA router and external networks.

Workaround

Create a low-preference, static, default route to the ABR.

Issue

(RN000243) If you use a virtual link to connect an OSPF area to the backbone router and if equal cost multipath (ECMP) routes from the area to the backbone router are set up, the OSPF virtual link does not form adjacencies with its neighboring routers. Because the adjacencies never reach the Full state, the OSPF area that is not directly connected to the backbone does not learn all routes in the autonomous system.

Workaround

Do not use ECMP routes when you use an OSPF virtual link.

Issue

(RN000242) If you use a virtual link to connect an OSPF area to the backbone router and if event logging of informative messages is enabled for OSPF, the following error message is generated when the virtual link receives a hello packet.

`Log entry <`*number*`> by event 12 at <`*date and time*`>: Informative(20)`

`Description: OSPF ERROR ignore area:0.0.0.0 mismatch from:` <*Sending Interface IP*> `on:` <*Receiving Interface IP*>

This message is displayed regardless of whether the hello packet is authenticated or not. However, the virtual link enters the Full state and the area that uses the virtual link to connect to the backbone learns routes correctly.

Workaround

Ignore the error message.

Issue

(RN000230) If you use a filtered search to search the OSPF link state database, and, when the switch displays the search results, you click **Refresh Table**, the switch displays the entire link state database.

Workaround

Perform the filtered search again.

Issue

(RN000168) If the switch learns equal cost multipath (ECMP) routes from multiple interfaces, and you use the Web Agent to search the routing table for:

- Routes learned from the higher numbered IP interface, the search results do not display the ECMP route.
- Routes learned from the lower numbered IP interface, the search results display both ECMP routes.

For example, if a route for destination 100.0.0.0 is learned from both the 50.0.0.1 interface and the 60.0.0.1 interface, and you use the Web Agent to search the routing table for:

- Routes learned from 60.0.0.1, the search results do not display the ECMP route.
- Routes learned from the 50.0.0.1 interface, the search results display both ECMP routes (through the 50.0.0.1 interface and through the 60.0.0.1 interface)

This is a display issue only. Although the switch learns and calculates the routes correctly, the routing table appears to have incorrect or incomplete routing entries.

Workaround

Use other criteria to search the routing table. You could search by protocol, for example, OSPF, or by a specific destination address, 100.0.0.0 in the preceding example.

Issue **(RN000170)** If you disable the advertisement setting for an OSPF summary, the switch does not retain the setting. When you reset the switch, it enables the advertisement setting.

Workaround Use the OSPF Summaries Web page to delete the summary if you do not want the IP network address that is defined in the summary to be advertised.

Issue **(RN000171)** The **no area <*area-id*> range <*ip-address*> <*mask*>** CLI command does not delete an OSPF summary.

Workaround Use the OSPF Summaries Web page in the Web Agent to delete OSPF summaries.

Issue **(RN000172)** When using the Web Agent, if you enter a 16-character MD5 key for an OSPF interface, the switch adds a symbol at the end of the key. This is a display issue only. MD5 authentication works correctly.

Workaround Enter a key of 15-characters or less.

Issue **(RN000173)** The CLI displays a range for the OSPF ***<poll-interval>*** parameter of 0-3600 seconds, even though the valid range for the ***<poll-interval>*** parameter is 1-3600.

Workaround Use the Web Agent or CLI to configure the OSPF <***poll-interval***> parameter to a value between 1 and 3600 seconds.

Point-to-Point Protocol (PPP) and Telnet

Issue **(RN000023)** A Telnet session to the serial port by means of PPP may time out during attempts to transfer large files, such as executable images, to a Trivial File Transfer Protocol (TFTP) server. However, this issue does not terminate the file transfer. An in-progress TFTP file transfer ends only after the file transfer is completed.

Workaround Not applicable.

Issue **(RN000024)** If you enter the **baud rate change** command more than once, the new baud rate may take effect before the current PPP connection is terminated. This results in the termination of the PPP connection.

Workaround Reestablish the PPP connection.

Fabric Mode 1 Port Mirroring

Issue

(RN000231) When you use Fabric mode 1 port mirroring to mirror a port range on an 80-series, 10/100 module, the last port in the range is not mirrored.

The switch allows you to enter a complete port range, for example ports 1 through 12 on a 24-port module. However only ports 1 through 11 are actually mirrored.

Workaround

Not applicable.

Issue

(RN000232) When you configure a port mirror on the Port Mirroring Configuration Web page and then click **Cancel**, you are not cancelled out of the page. Instead, the parameters for the port mirror are reset to their default values.

Workaround

None.

Issue

(RN000233) If you change the configuration of a port that is participating in a port mirror with a piggyback port, a topology change occurs. The port may lose connectivity for the duration of this topology change (approximately 30 to 45 seconds).

Workaround

If you need to change the configuration of a port and want the port to participate in a port mirror, change the port configuration before you enable the port mirror.

Issue

(RN000177) If you set up a port mirror on a 50-series switch that is operating in Fabric mode 1, and you monitor transmit traffic through a piggyback port, the frames that are received on the mirror port from the piggyback port are all tagged with a VLAN identification of ten. This issue occurs regardless of which VLAN number is bound to the source port. This issue is a display issue only. It has no effect on traffic patterns or flow.

Workaround

Not applicable.

80-Series QoS

Issue

(RN000259) If you use the Web Agent to configure an address forwarding table entry with a high or low priority setting on an 80-series module, the switch appears to accept the setting. Eighty-series modules support eight priority queues, which you can currently use

only the CLI to configure. Although 80-Series modules do not support high and low priority settings, the command is saved to the running configuration and the startup configuration files. When the switch runs the command, an error message is displayed in the event log file.

Workaround Use the CLI to set the QoS priority settings for 80-Series modules.

Radius

Issue **(RN000180)** If **cajun-service-type-required** is enabled on a switch that is running application software earlier than v5.3 and you upgrade the software to v5.3.*x*, the enabled setting is lost.

This issue occurs because the **set radius authentication cajun-service-type-required** command has been replaced by the new **set radius authentication switch-service-type-required** command. However, this issue occurs regardless of whether you used the CLI or Web Agent to enable the setting.

Workaround After upgrading the application software, reenable this setting. Use the RADIUS Web page or **set radius authentication switch-service-type-required** command.

RIP

Issue **(RN000264)** If you use the **network** command to enable RIP on an IP interface and inadvertently enter an IP interface that does not exist, the switch may enable RIP on an interface that does exist and on which RIP disabled.

Workaround Not applicable.

Issue **(RN000181)** When you use the Web Agent to configure the IP RIP key chain, the **Key Accept Time: minute** value is displayed incorrectly. The value displayed in the **Key Accept Time: minute** field will be equal to that of the value displayed in the **Key Accept Time: second** field. For example, if you wish the **Key Accept Time: hour** value to be 11, and the **Key Accept Time: minute** value to be 45, and the **Key Accept Time: second** value to be 10,

instead of the display 11:45:10, the switch displays 11:10:10. This is a display issue only, the parameter is added to the running configuration file (running.txt) correctly.

Workaround Use the CLI to configure key chains.

Issue **(RN000183)** When you configure the RIP key chain, if you set the **Key Accept Time** parameters to some time in the future, authenticated learning of RIP router interfaces does not begin at the time that you specify. Instead, authenticated learning of RIP router interfaces begins immediately.

This issue occurs in both the Web Agent and CLI.

Workaround Set the **Key Accept Time** parameters at the time that you want authenticated learning to begin. For example, if you want authenticated learning to begin at 13:07:25, set the **Key Accept Time** parameter at that time.

Use the Key Chain Web page or **key accept time <*key-id*>** CLI command to set the Key Accept Time parameter.

Issue **(RN000184)** When you configure the IP RIP key chain**,** and you set the **Key Name** field to a 16-character value, the string becomes corrupted in the running configuration (running.txt) file, and all communication to any network that is using the MD5 key chain is lost. Additionally, if the **Key Name** field is set to 16 characters, after a switch reset, the string is lost. Though the CLI online help specifies that the field should accept 16 characters, it does not.

This issue occurs in both the Web Agent and CLI.

Workaround Enter a value of 15 characters or less for the **Key Name** parameter.

Issue **(RN000185)** When you configure the key chain, if you enter a year later than 2009 or earlier than the year 1999 for the **Key Accept: Year** parameter, the display is corrupted, and the year is incorrectly displayed. Also, if you enter three characters instead of four for the **Key Accept: Year** parameter, authentication of RIP II MD5 does not work correctly.

This issue occurs in both the Web Agent and CLI.

Workaround Set the **Key Accept Year** to a year between 1999 and 2009. Do not enter a value of less than four characters.

Issue **(RN000186)** The CLI online help utility displays an invalid range value for the **Key Chain <*Key-ID*>** command. The help utility displays a maximum value of 2147483647 for the **<*Key-ID*>** parameter. The correct maximum value for the **<*Key-ID*>** parameter is 255.

Workaround Use a maximum value of 255 for the **Key Chain <*Key-ID*>** field when you configure it from the CLI.

Issue **(RN000187)** The online help utility for the CLI displays an invalid range value for the **timers basic <*basic*> <*invalid*>** command. The help utility displays a maximum value of -1 for the <***invalid***> parameter. If you set the variable to -1, a syntax error is generated.

Workaround Use the Web Agent or CLI to set the **<*invalid*>** parameter to a value from 1 to 65535.

Issue **(RN000188)** When you create or delete a duplicate RIP trusted neighbor, and then refresh the display, the most recently executed create or delete action is repeated. For example, if you create two duplicate RIP trusted neighbors, and then refresh the display, a third identical trusted neighbor is added to the trusted neighbor list. If the trusted neighbor list contains three duplicate RIP trusted neighbors, you delete one of the duplicates and then refresh the display, you will notice that another one of the duplicates has been deleted.

Workaround To configure RIP trusted neighbors, use the CLI.

Issue **(RN000135)** If you enable MD5 authentication on a RIP V2 interface and enter a nonexistent key chain in the **Auth Key/MD5 Key Chain** field, the Web Agent displays the following confirmation message:

```
RIP configuration was successful on interface
<xxx>
```

However, the MD5 authentication is not enabled.

Workaround Ensure that you enter an existing key chain in the **Auth Key/ MD5 Key Chain** field on the RIP Interfaces Web page.

Issue **(RN000033)** If you set an interface to use RIP V1/V2 and connect the interface to a RIP V1 interface, the RIP V1/V2 interface sends V2 packets to the V1 interface.

Workaround Set the interface to **RIP V1** if it is connected to a RIP V1 interface.

set module notes Command

Issue

(RN000266) If you use the **set module notes** command to add note text for a module and later delete the note text, the **show module** command, running configuration file, and Module Details Web page display corrupt data in the module notes.

Workaround

Use the Web Agent to delete module notes.

show ip route Command

Issue

(RN000189) If an interface name begins with a numeric character or contains a blank space, and you use the CLI to search the routing table for entries that are learned from the interface, the search results display only the default route, if one exists. If no default route exists, the CLI displays the following message:

```
There are no entries in the Route Table matching
search criteria.
```

The **show ip route <*interface-name*>** command searches the routing table for entries learned from a specific interface.

Workaround

Perform an unfiltered search by using the **show ip route** command or use the Web Agent to search for entries learned from a specific interface.

SNMP

Issue

(RN000244) If you use the **CajunOSPF.mib** to set the number of Equal Cost Multi Path (ECMP) routes on the switch, you are permitted to set it to an unlimited number. The switch supports a maximum of eight ECMP routes.

Workaround

The number of ECMP routes on the switch is set to eight, by default. It is highly recommended that you do not change this setting.

Issue

(RN000190) The promLogTableMaxSize object and promShutdownLogTableMaxSize object in the ProminetMIB.txt allow you to enter invalid values. If you enter invalid values, however, the settings revert to the default settings when you reset the switch.

Valid values for the promLogTableMaxSize object are 128, 512, 1024, and 2048.Valid values for the promShutdownLogTableMaxSize object are 16, 32, and 64.

Workaround

Use the General Event Management Web page or following CLI commands to set the size of the event log and shutdown log:

- **logging history size** for the event log
- **logging shutdown size** for the shutdown log

Issue

(RN000192) If you use the ProminetMIB.txt to configure SNMP trap logging, you are allowed to enable logging for certain event classes that you are prevented from enabling when you use the Web Agent and the CLI. When you reset the switch after saving the running configuration file to the startup configuration file, the invalid settings are lost, and the defaults are restored.

Workaround

To configure SNMP trap log settings, use the Web Agent or the CLI.

Issue

(RN000137) You cannot use SNMP to create an IPX static route. If you attempt to use SNMP to create an IPX static route, the IPX static route is not created.

Workaround

Use either the Web Agent or CLI to create an IPX static route.

Issue

(RN000138) The Avaya Multiservice switches currently support community strings of 29 characters or less. The Web Agent allows you to enter 30 characters, and the CLI allows you to enter 31 characters. However, the switch does not process PDUs with a community string of more than 29 characters.

CAUTION:
Avaya highly recommends that you set community strings to at least six characters and change all community strings to something other than public, their default setting. For more information on SNMP security and preventative actions that you can take, see http://support.avaya.com/security/2002-1/index.jhtml.

Workaround

Create community strings of 29 characters or less.

Issue **(RN000036)** If you attempt to remove the Public community string from the SNMP Community Management, save the running configuration to the startup configuration, and then reset the switch, the Public community string that you attempted to remove is not deleted.

Workaround Set the access for the Public community string to **None** and then save the running configuration to the startup configuration.

Spanning Tree

Issue **(RN000193)** If a port is disabled on an 80-Series gigabit module, the Spanning Tree Bridge Port Information Web page correctly indicates that the port is in the disabled state. If you disable Spanning Tree for the switch, and then enable it back in any of the Spanning Tree modes, the Port State for the disabled port incorrectly shows it to be in the Forwarding State. This is a display issue only, as the port is not actually in the Forwarding state.

Workaround Enable the port, then disable it again.

Static Routes

Issue **(RN000194)** If you use the **ip route-preference {local | rip | ospf-intra | ospf-inter | ospf-extra | static-hp | static-lp} <*value*>** CLI command to assign a preference to a route and enter an invalid number for the <*value*> parameter, the switch sets the preference to some other value.

Workaround Enter a valid preference value. Valid preference values range from 0 to 255.

Issue **(RN000139)** You cannot use the CLI to modify a static route. If you attempt to use the **ip route <*ip-address*> <*mask*> <*next-hop*> <*cost*> [{high | low}]** CLI command to modify a static route, the routing table will not reflect the change that you make.

Workaround Either use the Web Agent to modify the static route or, if you want to use the CLI, delete the static route and then create a new one with the new settings.

Summer Time Hours

Issue

(RN000261) When you use either the Web Agent or the CLI to adjust the system clock on the switch for a one-time change to the **Summer Time** or **Winter Time** hours, and the year that you enter for the one-time change is between the year 2000 and the year 2009, the switch truncates the number, and a year of 200 through 209 is saved to the running configuration and startup configuration files. If you enter a year greater than 2009 or less than 1999, it works properly.

Additionally, if you have made the one-time change, and then save the running-config.txt to the startup-config.txt file and then reset the switch, the **Summer Time** feature is disabled and you cannot reenable it until the erroneous number is removed from the file.

Workaround

When you configure Summer Time hours, use the recurring setting option. You can also manually reset the clock and NOT use the Summer Time feature. Additionally, if you are using the Simple Network Time Protocol (SNTP), you can configure the switch to appear as if it is in a different time zone by using the GMT Offset feature of SNTP.

If you have previously configured Summer Time hours for a one-time change, and need to remove the incorrect information from the running-config.txt file, use the following procedure to edit the file, and then copy it back to the switch.

1. Transfer the running-config.txt file to a TFTP server.
2. Use a text editor to remove the

   ```
   clock summer-time date May 30 20x 02:00 Jul
   2 20x 02:00 60
   ```

   line from the file where x is a number between 0 and 9
3. Transfer the edited file from the TFTP server back to the startupconfig.txt file on the switch.
4. When you reset the switch after you make this change, do not save the running-config.txt file to the startup-config.txt file, as your most recent change will be overwritten.

Supervisor Redundancy

Issue

(RN000198) If you use the Web Agent to configure CPU redundancy, you cannot delete the IP address that you previously configured for the CPU in slot 1 or slot 2 or default gateway. If you try to delete the IP addresses, or try to set them back to their initial default value of 0.0.0.0, the switch displays the following error message:

```
Invalid slot 1 console IP address specified
```

If this error message is displayed, you must click your browser's **Refresh** button for the CPU Redundancy Web page to display the originally configured IP addresses. The Web page is not automatically refreshed.

Workaround

To reconfigure CPU redundancy, use the CLI.

To assign an IP address to the slot 1 console, slot 2 console or the default gateway, use the **cpu_redundancy console <*slot1* | *slot2* | *default-gateway*> <*ip-addr*>** CLI command:

To remove or reverse this configuration, use the **no cpu_redundancy console <*slot 1* | *slot2* | *default-gateway*>** CLI command.

Switch Controller

Issue

(RN000044) When the switch controller in a P880 switch with a 50-series supervisor fails, the switch provides no error message that specifically states that the controller has failed. If the switch has a redundant switch controller, the switch resumes normal operation after the primary switch controller fails over to the redundant switch controller. If the switch does not have a redundant switch controller, you will need a replacement switch controller.

Workaround

Not applicable.

System Clock

Issue

(RN000222) If you enter a specific time for recurring summer-time hours, the system changes the time at midnight, regardless of the time that you enter.

This issue may result in log entries not listing the correct time.

Workaround

Manually change the current time on the system clock.

Upgrading the Application Software

Issue

(RN000239) If you attempt to view the FEPROM Web page while you are downloading application software to the switch, the Web Agent stops responding.

Workaround

Use the **show flash** CLI command to check the status of the download.

Issue

(RN000237) If while upgrading the application software, you use the **show cpu status** CLI command or the CPU Status Web page to view the redundancy status of the supervisor modules, the switch displays the following error message:

```
Erase timeout at addr 0xFFE30000, status=0x7AWrite
failed, addr 0xFFE00001 not erased.

ERROR[Failed write...transfer cancelled.] file
'fagin.bin', TFTP server 199.93.238.49
```

Workaround

Do not attempt to view the redundancy status of the supervisor modules during an upgrade of the application software.

VLANs

Issue

(RN000249) If a 10-Gigabit port is bound to a VLAN, but nothing is physically connected to the port, the following Web page and CLI display the 10-Gigabit port number incorrectly:

- The VLAN Switch Ports Web page displays ***<slot N>-<slot N+1>*.8** instead of ***<slot N>-<slot N+1>*.1**.
- The **show vlan** CLI command displays ***<slot N>-<slot N+1>*/8** instead of ***<slot N>-<slot N+1>*/1**.

For example, if 10-Gigabit port 7-8.1 is bound to VLAN 10, but nothing is connected to the port:

- The VLAN Switch Ports Web page displays **7-8.8** instead of **7-8.1**.
- The **show vlan 10** command displays **7-8/8** instead of **7-8/1**.

Workaround

Not applicable.

Issue

(RN000212) If you change the **VLAN binding** setting from **Bind to All** to **Bind to Receive** for a trunk port, a loss of connectivity may occur in the network.

Workaround

After changing the **VLAN binding** setting from **Bind to All** to **Bind to Receive**, delete the AFT entries for the VLANs that are bound to the trunk port.

To delete the AFT entries for a specific VLAN:

1. Open the VLAN Configuration Web page (**L2Switching** > **VLANs** > **Configuration**).
2. In the **Table Index** field, click the index number that is associated with the VLAN for which you want to delete AFT entries.

   The Address Table Instance Web page is displayed.
3. Click **Delete All Learned Entries**.

Issue

(RN000226) If you want a port to use both the VLAN auto-learning feature and VLAN binding type of **Bind to Received** or **Bind to All**, the port may not be automatically added to the VLAN if you do not set the VLAN binding type first.

Workaround

If you want a port to use both the VLAN auto-learning feature and VLAN binding type of **Bind to Received** or **Bind to All**, set the VLAN binding type *before* you set Auto-learn to **Enable**.

VRRP

Issue

(RN000258) If you are running VRRP, the master router may become the backup, and the backup router may become the master if you change one of the following settings:

- Auth-Key
- Address Owner Override
- Preempt mode

Workaround

Make changes to the VRRP configuration during a maintenance window.

Issue

(RN000235) The switch does not support both Proxy ARP and VRRP on the same IP interface.

Workaround

Not applicable.

Issue

(RN000213) When you change or edit any parameter on an interface that has VRRP enabled, the event log displays the following informational message:

```
269 5 02-Mar-28 11:11:10 Informative(20) Status
VRRP:FSM:VrrpProcEvent on interface

VRRP1 VRID 1: event STARTUP state MASTER --> MASTER
```

Workaround

Ignore the message.

Issue

(RN000082) You cannot change the IP address of a virtual router by using the CLI.

Workaround

Use the Web Agent to change the IP address of a virtual router.

Technical Support

To contact Avaya Technical Support:

*** Note:** These are new phone numbers as of October 1, 2000.

- From the United States:

  1-800-237-0016

- From North America:

  1-800-242-2121

- Outside North America:

  Contact your distributor

Documentation Feedback

If you have comments about the technical accuracy or general quality of this document, please send us an e-mail message at:

LSG-CTechPubs@avayactc.com

Please cite the document title, part number, and page reference, if appropriate.