



Ethernet Routing Switch 8800/8600

Readme

Software Release 7.2.1.1



Table of Content

Software Release 7.2.1.1	3
Important Notices.....	3
Supported Hardware for Release 7.2.1.1	3
Upgrade Notes.....	3
Important Operational Notes.....	5
File Names for This Release	5
Version of Previous Release.....	7
Compatibility	7
Changes in This Release	7
New Features in This Release.....	7
Old Features Removed From This Release.....	8
Problems Resolved in This Release	8
Outstanding Issues.....	13
Known Limitations	14
Documentation Corrections.....	15



Software Release 7.2.1.1

Release Date: June 19, 2013

Purpose: Software maintenance release to address customer reported issues.

Important Notices

Supported Hardware for Release 7.2.1.1

Refer to the ERS 8800 Release Notes for a complete list of supported hardware **Release Notes — Software Release 7.2.0.0 Avaya Ethernet Routing Switch 8800/8600 (NN46205-402, 09.02)** available at <https://support.avaya.com>

Upgrade Notes

This section describes the Ethernet Routing Switch 8800 Software Release 7.2.1.1 software files.

Before you upgrade, it is recommended to verify the MD5 signature for each new file to be used.

For upgrade procedures, see **Upgrades — Software Release 7.2.0.0 Avaya Ethernet Routing Switch 8800/8600 (NN46205-400, 07.2.0.0)** available at <https://support.avaya.com>

Important Upgrade Note for systems using NLB configuration:

Due to a software issue present in previous releases, the system saves the “ip arp static-mcast” related configuration to the config file in an incorrect way as listed below. The parts shown in red are incorrect and should not have been saved. Loss of configuration can result upon system boot up when booting from a config file with this error in it. The reason is that the config-loading operation, that happens upon system boot up, stops prematurely and configuration data that falls beyond the portion below does not get applied.

```
router vrf <vrf name>
ip arp static-mcast <ip add> <mcast Mac@> vid <vlan id> <ports>
exit
```

If these extra fields exist in a currently saved config file, these “exit” lines must be manually removed from the saved config by getting the file from the ERS 8x00 to a workstation, editing the file (using Notepad) and then downloading the corrected configuration file back onto the ERS 8x00 device. It is strongly recommended that a backup copy of the clean config file be saved off on the flash to use in case of emergency. Note that any time a save config is performed, the incorrect fields will get saved again and will need correcting.

This problem has been fixed in this release. However if upgrading a system from an older release to 7.2.1.1, care needs to be taken to ensure the upgrade is performed using a config file that is free from these errors.



Important Upgrade Note for systems with 8895 CPUs and Dual Input Power Supplies (8005DI 1170W/1492W 120/240V Dual AC-input Power Supply):

Background:

- An I2C bus hung state can result in systems with 8895 CPUs and 8005 Dual Input (DI) Power supplies running pre-7.1.5.2 or pre-7.2.1.1 release. The system will continue to function properly even with the I2C bus in a hung state. However when system is rebooted, the reboot will fail due to “Chassis SEEPROM read error” and device will remain in a software hung state and potentially become inaccessible.
- This issue is fixed in release 7.1.5.2 and higher software release versions and in 7.2.1.1 and higher software release versions.
- While normal operation is not impacted by the I2C bus being in a hung state, switch initialization (reboot) /S impacted and additional steps need to be taken as part of the upgrade to ensure the switch recovers properly. Failing to follow the additional steps (see ‘special procedure’) below will result in a failed upgrade and the device will remain in a software hung state upon reboot and potentially become inaccessible.

Monitoring:

- The device can be monitored periodically using the “show sys info” (ACLI: show sys-info) command. Note: Only run this command on the Primary CPU.
- If the Fan temperature reading shows “0”, it is indicative of a hung I2C.
- The system will continue to function properly without issue even with the I2C bus in a hung state and will only pose issues upon reboot & upgrades. See the ‘special procedure’ in the ‘Upgrade Instructions’ section below on how to deal with such a situation.

Upgrade Steps:

- Ensure someone is physically present near the device, in case a power cycle is needed to recover the system.
- Download the 7.2.1.1 or higher software release onto the device and get all the pre-sets done as you would in a normal upgrade. However DO NOT issue the reset or reboot command. - Run the “show sys info” command. If the Fan temperature reading shows “0”, it is indicative of a hung I2C. If this is the case, you need to follow special procedure below:
 1. Once all pre-work is complete, power the chassis off
 2. Remove **ALL** power supplies from chassis and hold them out for a minimum of **30** seconds. This allows the capacitors in the supply to discharge and ensure the supply has gone completely cold electrically. This step ensures the I2C bus will start clean on power-up.
 3. Reinsert all the necessary power supplies into the chassis securely.
 4. Power on the chassis
 - a) Make sure you have at least one PC with working console access to CPU (start with CPU in slot 5 if only one console terminal available).
 - b) (for chassis with dual CPUs): pull out CPU in slot 6
 - c) Power up the chassis power supplies (all) and monitor console output
 - d) hit <RETURN> repeatedly and stop the booting process at monitor mode
 - e) at monitor mode, run command: ‘ boot /flash/p80b<new_sw_load>.img’
 - Chassis will reboot again and this time do not hit any ENTER but just watch on console as CPU powers-up normally
 - If chassis has only one CPU, then go to step 5. If chassis has two CPUs then go to next step, f.
 - f) connect console and terminal to console port on CPU-6
 - g) insert spare CPU in slot-6
 - h) repeat steps d – e on CPU-6 console.
 5. The chassis should boot successfully with the 7.2.1.1 or higher software that it was preconfigured to load.
- If fan temperature readings were normal, you do not need to do the above procedure. You can simply complete the upgrade using the normally outlined upgrade procedures.
- The chassis should boot successfully with the 7.2.1.1 or higher software that it was preconfigured to load.
- If for any reason the chassis is hung during software initialization during this upgrade boot, you can recover it using the above noted special procedure.



Important Operational Notes

- Do not use TFTP to transfer large files to/from the device. There is a possibility of causing stack overflow and potential crash of the system FTP should be used instead.
- Out of band management network addressing scheme should be kept completely disjoint from in-band network addressing, ensuring that no overlap results. Any overlap between the management network and in-band network can result in network loops and severely affect network performance. A configuration with overlap between management and user address schemes is not supported.
- The “verbose mode” feature of CLI/ACLI is to be used to generate a verbose mode configuration for the sole purpose of viewing default configuration options. Sourcing or loading of such a “verbose” config file is not supported. Only a non-verbose configuration should be used for device operation.
- Configuration of non-IST inter-switch links running parallel to an inter-switch IST is not a valid configuration and can result in network loops. This kind of configuration should never be done.
- It is important to make sure that a system upgrade does not result in a condition where the system boots with mismatched boot-monitor image (“b” or “be” image) and run-time image (“a”, “ae” or “m” image). This can result in incorrect system behaviour. Please follow correct procedure to boot with the boot-monitor image
- Special care must be taken when upgrading dual-CP systems running in hot-standby HA mode. The two CPs need to be booted in a specific order within a very specific time gap between them. Booting outside of these guidelines can result in upgrade failure. Further, upgrade of IO cards will not take place if this procedure is not followed correctly resulting in software release mismatch between CP and IO cards resulting in malfunction of the system.
- Show commands such as “show sys info” that access system level information or SNMP gets/walks that access system MIB should never be run from/against the Backup CPU.
- From the time the link establishment process begins, 10G links on 8812XL cards may take a minute or so to become operational.
- SMLT / IST ports enter a special “hold-down” state transition immediately after reboot, as the system comes up and until it reaches full operational state. As such any dynamic configuration changes or movement of SMLT/IST ports should not be done until system is completely up and all ports have reached their normal run-time states.

File Names for This Release

Module or file Type	Description	File name	Size in bytes
Tar file of all software	Deliverables (includes images that also contain encryption software)	pr86_7211.tar.gz	67604489
Ethernet Routing Switch images			
Boot monitor image for 8692 SF/CPU	8692 CPU and switch fabric firmware	p80b7211.img	1192681
Boot monitor image for 8895 SF/CPU	8895 CPU and switch fabric firmware	p80be7211.img	1263506
Run-time image for 8692 SF/CPU	Run-time image for 8692 SF/CPU	p80a7211.img	16515226
Run-time image for 8895 SF/CPU	Run-time image for 8895 SF/CPU	p80ae7211.img	15491133



Run-time image for R modules	Run-time image for R modules	p80j7211.dld	1905700
Run-time image for RS modules	Run-time image for RS and 8800 modules	p80k7211.dld	1971972
Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz)	Image for the SuperMezz Card	p80m7211.img	16594764
3DES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c7211.des	56124
3DES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce7211.des	51972
AES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c7211.aes (this image includes the 3DES image)	27436
AES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce7211.aes (this image includes the 3DES image)	25156
MIB	MIB files	p80a7211.mib	5421612
MIB (zip file)	Zip file containing MIBs	p80a7211.mib.zip	852275
MD5 checksum File	md5 checksums of all Release 7.2.1.1 software files	p80a7211.md5	1234
FOQ for R modules	Feedback output queuing FPGA firmware	foq267.xsvf	5320469
BMC for R modules	BAP memory controller FPGA firmware	bmc776.xsvf	2640266
DPC for R modules	Dual port Controller FPGA firmware	dpc194.xsvf	2642001
PIM8630GBR	Programmable I/O module FPGA firmware; for the 8630GBR only	PI_769.xsvf	2284578
Firmware for RS modules	Contains FOQ, BMC, DPC, mirroring, and loopback images	rs_dpm_fpga.bin	4538368



PIM images for RS modules	PIM FPGA firmware required for 8612XLRS module only	pim8612XLRS.bin	60183
	PIM FPGA firmware required for 8634XGRS module only	pim8634XGRS.bin	78173
	PIM FPGA firmware required for 8648GBRS module only	pim8648GBRS.bin	79891
	PIM FPGA firmware required for 8648GTRS module only	pim8648GTRS.bin	54441
Trace files			
MPLS trace file	Trace file for MPLS. This is auto generated and appears on the PCMCIA after upgrade.	nbpdtrc.lo0	
EDM Help files			
EDM help files	Help files for EDM GUI	ers8000v7210_HEL P_EDM.gzip.zip	3852274
ERS 8000 EDM plug-in for COM			
EDM plug-in for COM	EDM plug-in for COM	ers8000v7.2.1.1.zip	5324405
Copyright	File containing copyright information	ERS8k.7.2.1.1_Copyright.pdf	266541

Version of Previous Release

Software Version **7.2.0.2**

Compatibility

This software release is managed with Enterprise Device Manager (EDM).

Note: Updated EDM Help files and a new EDM plug-in version (v7.2.1.1) for COM are included with this release. Those must also be updated as part of upgrading to release 7.2.1.1.

Changes in This Release

New Features in This Release

- TACACS authentication was not supported for EDM. It continued to use the locally set password and did not consult the TACACS+ Server. The ability to use TACACS authentication for EDM connections has now been added. If both TACACS and Radius authentication are enabled, system tries the TACACS authentication first. If TACACS's server is unreachable, Radius authentication is attempted. If Radius also is unreachable, system falls back to local authentication [wi01086151]
- WLAN SPLIT PLANE enable option is now available starting from this release [wi01099014]
- In the output display of 'show ip ospf lsdb' command, Summary and AS-Summary LSAs should also display the metric value. Currently, the metric information was displayed only for AS-External LSAs. Now the metric is



displayed for Summary LSAs and AS-Summary LSAs as well. This functionality is implemented for CLI and ACLI [wi00988777]

- The command “show isis spbm ip-unicast-fib all spbm-nh-as-mac” does not display VRF(L3VSN) fib entries. Added that capability [wi01044371]
- Unable to access the ERS8xxx via EDM using IE8 after applying Windows-XP security patch. ERS 8xxx supplied 512 bit certificates for SSL connections. Now ERS 8xxx also supports 1024 bit certificates [wi01079126 / wi01061189]

Old Features Removed From This Release

None

Problems Resolved in This Release

- For a dual-stack-enabled (IPv4 and IPv6) VLAN with RSMLT edge support enabled, upon an HA failover, the RSMLT timers get reset and start from initial value instead of continuing from current value. This issue has been fixed. Also fixed the problem where RSMLT holdup timer not counting in edge support scenario and peer MAC is made local permanently. [wi01035599]
- ERS8800 Radius CLI Profile now works with SSH [wi01041501]
- ACL problem: NextHopReachableRedirect - When the redirection destination is not reachable the filter continues to redirect instead of dropping. This is now fixed [wi01041516]
- Inconsistent ACL/ACE behaviour is observed on 8648GTRS module [wi01041506]
- show fulltech can be disallowed if SSH session is stranded during execution [wi01041491]
- ACLI: Access-policy-SNMPv3 configuration disappears after re-boot [wi01041523]
- Syslog message were sent using LAST created CLIP Interface IP address instead of the first configured CLIP address [wi01041521]
- ERS8800 with CP card 8895SF/CPU registered three consecutive core dumps. Device running on s/w 7.1.3.0. SOP init problem [wi01041511]
- A possible crash could happen when running "show fulltech" during "show ip count-summary" section. This crash has been fixed [wi01042922]
- In-service replacement of an IST peer chassis could result in a condition where the non-replaced chassis continues to use the wrong (older chassis's) MAC address in the SLPP packets. This issue has been fixed. Also fixed a potential bad memory access situation in the MAC delete by port function [wi01037846]
- Syslog produces incorrect host numbers when enabling or disabling the service. This was caused because the CLI function did not pass correct host id to function handling the "no syslog host <id> enable" command. This has been fixed. [wi01038842]
- Device sends LACP with irregular timer on average every 15 sec when configured for long timer of 30 seconds. This was caused by two independent tasks generating the LACP PDUs. Added synchronization between these tasks to ensure a task does not send the PDU if other task has sent it for the said interval [wi01044664]
- The ACLI does not have a command for removing an OSPF stub area like there is in the CLI. This was resolved by adding a stub option under the NO form of the “area” command. [wi01075445]



- The VRF IP static multicast MAC ARP entries were causing two issues: an invalid error message and incorrect behavior. The invalid message appeared because a deleted entry was still appearing in the show command. A comparison check was added to avoid this issue. The incorrect behavior was that multiple exit entries were appearing in the configuration file. After an exit, the flag was set properly to prevent this from happening [wi01075452]
- After an OSPF FULL adjacency is established, if the last entry in the Link State Database has a nonzero forwarding entry, then any new external routes received in an LSUPDATE will be injected in to the routing table with that forwarding entry as their next hop. This issue is fixed. [wi01079116 / wi01059293 / wi01044656]
- When the switch received a BGP Update message whose path attribute was greater than 260 bytes, it corrupted the memory area beyond 260 bytes and caused a crash. This crash has been fixed [wi01079122]
- After a reboot, the routed packets that should have routed out the VRF (Default Gateway) were routed back and forth between the two office ERS 8600s. This issue was fixed by adding an IP default route and checking the source VRF id [wi01079127]
- If you insert a line card and the default route is on VRF 1, the following messages appear: "COP-SW ERROR" rcdAddEcmpDefaultRoute: Arp rcdRadixLookup failed" and "ercdProcArpRecMsg: Failed to Add Ecmp Default Route." A new function fixes this issue by installing the ARP records correctly [wi01079131]
- The ERS 8800 is unable to apply a route policy to the redistributed routes from one VRF to another. Setting the redistribute table route map fixed this issue. [wi01035373]
- Ping and traceroute are not functional when sourced from one VRF to a device of another VRF. This issue is now fixed. [wi01035348]
- Qualys scan related robustness fixes have been made [wi01078774]
- Multiple PIMSM crashes in memory instrumentation image [wi01079638]
- Support for topology discovery for VSP4k [wi01079637]
- When the last port of STG1 is removed, all ports added back to STG1 after config save [wi01079634]
- The error message "Stop polling DI Power Supply due to excessive i2c error; fan temp=0" would be seen in system running 8895 and DI Power supplies. In systems with 8895 CPs, there is a potential for the I2C bus to get stalled. This could result in peripheral devices on the I2C bus becoming unreachable from the CPU [wi01079445 / wi01058836] **Special care should be taken when upgrading a system running pre-7.2.1.1 software and having 8895s and Dual Input power supplies. Please see the Upgrade note for details.**
- In some cases an incorrect MAC address value would be displayed in the SLPP log entries related to loop detection. This has now been corrected [wi01076837]
- A metric set through the route policy was not applied to the IP shortcuts case. This affects both the GRT (TLV 135) and SPBM IPVPN (TLV 184). Added support to allow the metric to be manipulated through route policy [wi01083344 / wi00853802]
- ISIS adjacencies up-time drifts by 5-6 seconds per minute. This has now been fixed [wi01083344 / wi01053163]
- The nick-name option in show isis spbm multicast-fib sometimes displays the system name for the wrong BEB. A code change to compare three bytes of the nick-name instead of two fixes this issue. [wi01058539]



- A potential for system reset resulting from SNMP trapd task exception during a memory freeing operation has been corrected [wi01082834 / wi01081346]
- The default trace level for Infinity module was set to VERY_TERSE instead of NO_DISPLAY [wi01085280 / wi01039790]
- No ability to configure Inter-VRF route redistribution via EDM. Added support to allow this configuration [wi01085689 / wi01075385 / wi00993971 / wi01043317]
- In ACLI mode, configuring an ospf area range summary for the default area and saving the configuration results in an invalid saved configuration as the context was not set correctly to "router ospf" [wi01085685 wi01046124 / wi01066179] - Upgrade Warning: If upgrading a system that uses this configuration, the saved config from older release will need to be manually corrected and used for the upgrade to avoid configuration loss upon reboot.
- When you reset two BEBs that are IST peers, some of the L3 streams do not recover because the individual streams are using different ISIDs after the reboot. This issue was fixed by clearing the old LSP ISID information. [wi01085529]
- When you create a black hole default route and reboot, it causes a system crash because RSMLT tries to add a temporary default route first. This issue was fixed by checking to see if the route being added is a black hole. [wi01086324 (port from wi1082670)]
- Running SPB while configuring IGMP Snoop on VLAN caused a crash. This crash has been fixed. [wi01086486 / wi01086641]
- Upon user-initiated device reset from CLI, the logging operation waiting to flush the logs before proceeding to reset can cause a deadlock and force a watchdog reset. This issue is fixed. [wi01046166]
- Fixed an issue with IPFIX where the data would be dropped if the active timeout and the expiry intervals happened to be exact multiples of each other [wi01003993]
- BGP message exchange between an ERS 8xxx peering with a host running Linux OS can result in a crash of the switch. Some Linux implementations are incorrectly sending update message with the "AS path attribute length extension" bit set even when the length is less than or equal to 255 bytes. The ERS 8xxx now handles this condition by adjusting the bit, if necessary, by checking the actual length of the update message. [wi01077643 / wi01018835]
- ERS 8xxx switch can experience a CPU switchover due to memory corruption. Robustness fixes have been made in the areas of BGP, RADIUS, Route Redistribution, Secure Copy (scp) and Ethernet [wi01077643 / wi01045784, wi01049417]
- IGMP V3 related memory corruption has been fixed [wi01077643 / wi01060672]
- IPFix configuration is only allowed in the default VRF context. IPFIX packets from certain slots could be wrongly interpreted as coming from/destined to a wrong VRF resulting in IPFIX packets not being sent out from certain slots. This issue is fixed by forcing all IPFix collection packets to VRF 0 [wi01077643 / wi01062810]
- COP Record corruption and subsequent crash could result due to a race condition in the handling of FTN and IP routes in the RTM. This issue has been fixed [wi01065076]
- Https related crash can result from running out of file descriptors. This issue has been fixed [wi01065671]
- A null pointer crash in the BGP area resulting from processing of a state change event has been fixed [wi01065676]



- CP-Limit feature does not work on ports on slots 8 through 10 on release 7.1.x.x (older than 7.1.5.1) and release 7.2.0.2. Fixed the CP-limit feature to shut the offending ports down for slots 8, 9 and 10 [wi01065450]
- Crash seen on ERS 8xxx when performing VPFM discoveries against ERS 8x00 seeds has been fixed. This crash can occur when the ERS 8xxx has both AES and 3DES encryption modules loaded and VPFM includes both SSH and SNMP credential configuration [wi01066759]
- Fixed CPP FBUF loss that would occur due to presence of incoming traffic during bootup. This was visible in HA failover scenario with IPFIX enabled.
- Redistribution of black hole static route results in a log error message: "SW ERROR Route 10.166.0.0/255.255.0.0 index 0 has null next hop pointer". There was no impact to traffic flowing in the network. The problem was fixed by checking for black holes. [wi01085998]
- Security enhancement related to IPFix feature has been added [wi01089731 / wi01085229]
- SSH/SSL related memory leak and memory corruption has been fixed. [wi01086607 / wi01086611]
- ARP delete operation has been fixed to ensure proper removal of ARP entries from hardware [wi01089721 / wi01089705 / wi01088321]
- SFP info does not show on reboot with admin disabled ports. SFP information was specifically bypassed for disabled ports to reduce polling overhead; ensured reads happen on all transitions [wi01088690]
- When an ARP request is sent out for the NLB virtual IP address (VIP) in the "nlb multicast" mode, all the hosts within that cluster reply back with a multicast MAC address. The ERS 8600 was adding the MLT/port on which the reply is received. On one of the SMLT peers, this is going to be the IST. Henceforth, any traffic destined to the VIP from that ERS will be sent on the IST and that will be dropped by the peer. To fix this issue and prevent the peer from dropping traffic, the ERS now learns the port/MLT of the SRC Mac of the host sending back the ARP reply. [wi01064178]
- IGMP control traffic received on a UNI for a L2VSN is sent to the CPU even though IGMP-SNOOPING or Multicast-Routing is NOT enabled on the VLAN. This behavior is incorrect. IGMP control traffic should only be sent to the CP if IGMP-SNOOPING or Multicast-Routing is enabled on the VLAN. This problem was a 7.2 issue only and it has been fixed. [wi01094501]
- When the MAC is learned on an NNI VLAN and that VLAN does not have an ISID configured on it, it was being incorrectly marked as NNI when the information was synced to the IST peer. [wi01088133]
- When the ISIS system-id changes on the remote end without giving the local end a chance to declare the adjacency down, the ISIS on the local node is kept in the INIT state and the remote node goes down. Because the local node still had the adjacency up when the hello packet from the remote came in with the new system - id, it was still sending the hello packets with the old system-id kept in its adjacency pointer, which was rejected by the other side. This issue was fixed by comparing the remote system -id in the adjacency pointer with the remote system -id in the header every time a hello message is received, and resetting the adjacency if it does not match [wi01094156]
- Under certain conditions, an ERS 8600/8800 device that has IPFix enabled could see the Line Card (COP) CPU utilization spike to 100%, affecting other protocol operations or resulting in network instability/slowness. That issue has been fixed. An issue related to the accuracy of the data reported by IPFix has also been corrected [wi01061665]



- Memory robustness improvements (fix memory leaks resulting from "show fulltech" command [wi01096320], from performing SNMP operation on rcPrFilterAceProtoShowTable MIB have been fixed [wi01096322], from HA sync after creating filters for both ipv4 and ipv6 [wi01096332] have been fixed)
- Robustness improvement related to TACACS login [wi01097042]
- The mlt id argument noted on the ip arp static-mcast command does not load when loading config after boot because it is not accepted by the syntax. This is specific to the VRF based command. The GRT command works correctly
vlan static-mcastmac 4007 03:bf:0e:32:07:fc mlt 1,2,21,22
ip arp static-mcast 14.50.7.252 03:bf:0e:32:07:fc vid 4007 1,2,21,22
The command is also putting 8 extra "exits" in the config file resulting in the rest of the config to not get loaded. This causes partial config loss on reboot [wi01098181 / wi01095559]
- When configuring RIP to redistribute ISIS routes into RIP using a route-map, the configuration change were not getting saved into the config file [wi01098085]
- The MAC address record for the ISIS SYS ID should have the RTR Bit set when IPSC is enabled. It was found that this bit was turned off even though IPSC is enabled, causing all the IPSC pkts to routed thru the CP [wi01099642]

- Queries to unsupported ISIS MIBs could result in a device crash. The following MIBs are unsupported and if polled directly or indirectly, can result in a crash:

<u>MIB</u>	<u>OID</u>
isisSummAddrTable	1.3.6.1.3.37.1.1.4
isisSAAdjAreaAddrTable	1.3.6.1.3.37.1.6.2
isisSAAdjIPAddrTable	1.3.6.1.3.37.1.6.3
isisSAAdjProtSuppTable	1.3.6.1.3.37.1.6.4

This has now been fixed [wi01100357]

- The browse button for IP redistribution screen in VRF context does not show all VRFs. This problem has now been fixed. [wi01101725]
- Excessive flooding can result from NNI outgoing interface being incorrectly programmed for NNI traffic. [wi01091954]
- OSPF adjacency between access switches does not come up over L2VSN when IGMP snooping is enabled on c-vlan on BEB. This problem is now fixed. [wi01099485]
- Duplicated L2 Mcast traffic can result after reboot of core switches one by one. This is a result of a timing race condition. Issue is now fixed [wi01102946]
- An SPB BEB node can crash with the following signatures: ip_rdeleteVrf: orec is NULL! RT:%ld.%ld.%ld.%ld or plsbDeleteBmac: routeTree empty, non-zero count. This can happen in the following scenario:

Alternative route feature is disabled in the GRT or VRF under consideration

There is a network/route that is learned via some routing protocol other than ISIS (RIP, OSPF, BGP, static, etc)
The same route is learned via SPB/ISIS. SPB will add that route to the RTM, and if ISIS_SPBM route is more preferred, will replace the existing route

To replace the route, the RTM first deletes that existing "Best Route"

This triggers a notification to other protocols to submit their "Best" Routes



This trigger to ISIS/SPB, will cause ISIS/SPB to call RTM to add the same route again in a recursive manner, which can cause corruption of the RTM data structures.
The crash can occur, when next time the BEB tries to remove that route.

This has been fixed [wi01107370]

- When an IST link is physically removed, the IST does not recover if the link is left down for more than 8 minutes. [wi01055421]
- Traffic does not pass between a server with Intel 10G NIC connected to 8812 after the server is rebooted [wi01051880]
- Port does not initialize correctly with a 100Base FX SFP and the following message is reported: "KHI Warning Port <number> is experiencing Packet Errors, FCS Errors" [wi01066167]
- Under high levels of multicast traffic, the 8630GBR experienced resets due to false packet memory refresh errors [wi01064252]
- A user with read-only privilege is not able to launch a non-Global VRF Context View using COM and EDM Plug-in. EDM Plug-in v7.2.1.1 addresses the above issue. A read-only user can now launch a non-Global VRF Context View from COM Application [wi01071995]

Outstanding Issues

- When directed broadcast is enabled on a VLAN with IP – the expected behavior is that all port members of the VLAN receive copies of the packets addressed to the subnet broadcast address of the VLAN IP interface. When the VLAN is also a L2VSN – the expectation is that copies of the packet are sent out of all local UNI port members of the L2VSN as well as UNI port members of the L2VSN connected to other BEBs across the SPB Network. These packets are currently sent only to local UNI port members on the BEB that is routing the directed broadcast packets.
- With the IPv4 feature “ip more-specific-non-local-route” enabled, when a host route already learned on a local network is also learned on a non-local network, subsequent withdrawal of the non-locally learned route leaves the host on the local network unreachable due to a problem with relearning the original ARP. This is an advanced feature, not applicable in most networks. It should not be enabled without a specific need. This issue will be fixed in an upcoming release. A fix for this is available in 7.1.5.3 release [wi01088319, wi01087883]
- 8812XL IO modules utilizing copper direct attach cables (DACs), Avaya part numbers AA1403018-AA1403020 inclusive, may occasionally miss an interrupt when the DAC is removed from the module. This may cause the link state to remain up even though no transceiver is present. The front panel light for the port may remain lit and the CLI/ACLI/EDM may indicate the link is in its active state. To avoid this condition, users should administratively disable ports on which copper DACs are deployed prior to removal. There are no known issues with DAC insertions. In all cases, link will be in the correct state on insertion. If the removal condition is inadvertently encountered, any action which would result in a link state change to any other port on the slot will clear the condition. If using DACs for inter-switch links (instead of fiber), it is possible that switch would not detect any link failure. Mechanisms such as VLACP should be used for inter-switch links using DACs to ensure proper link failure detection. [wi01106552]
- A tMainTask loop condition could result from TFTP task stack overflow resulting in a watchdog triggered reset. Avoid use of TFTP to transfer files to/from the device. FTP should instead be used [wi01062093]
- 10G Link STP Physical Path Cost is not recalculated to the expected value of 1. If a 10G port is removed and re-added to the STG, the path cost is being set to a value of 100 which can cause other ports in the network to block and values are incorrect throughout the network. Workaround is to manually configure the path cost to be 1 [wi01087348]



- In ACLI mode, the syntax to enable and configure RIP parameters on multiple Ethernet interfaces is as follows. When the list of ports were supplied, then the below settings did not apply to all of the ethernet ports.

```
interface GigabitEthernet 3/2,3/3  
  brouter port 3/2 vlan 3902 subnet 192.168.255.6/30  
  brouter port 3/3 vlan 3901 subnet 192.168.255.2/30  
  ip rip enable  
  ip rip receive version rip2  
  ip rip send version rip2  
exit
```

The highlighted commands were not applied to all the ports when the list of ports were provided as a parameter to the command. Workaround is to do run these commands individually on each port. CLI does not have this problem.

Known Limitations

- When the number of files stored on PCMCIA card or internal flash exceeds 100 or the free space remaining on the PCMCIA or internal flash gets low, the I/O modules may experience a delay in coming online during an upgrade or downgrade. Also the system may experience general slowness at run time. This can be easily avoided by clearing unused files from those devices and ensuring there is sufficient free space on PCMCIA card or internal flash to avoid this condition [wi01100333]
- During the boot up sequence, on an ERS 8600 system with a single 8692 CP, intermittently, the following message may print to the console/logs. This message can be ignored. It does not affect the functionality of the system. This message is not seen on systems with dual CP cards or on 8800 systems having either single or dual 8895 CPs.
HW ERROR Code=0x37 iccSendEcho: sendto 127.0.0.5 status -1 expected 32
- Link Flap and automatic port shutdown may result for the ports connected to 8812 ports when the system is booting up. The link flap duration is about 20 seconds and is directly related to the number of ports in use. The duration is lowered with the fewer number of ports used. There are two workaround options:
 1. Increase the link flap detect interval to a number greater than the duration of the link flap seen.
 2. Disable link flap detect auto port down (not recommended)
- When the operating speed/duplex on a port is changed, a momentary traffic loss is seen on other ports in the lane. The reason is that every time a port speed changes, lane-specific shapers need to be reconfigured. This is also the case if Egress Queue Sets are being reconfigured or port shapers are being configured.
- Please also refer to the Known Limitations section of the Release Notes for Ethernet Routing Switch 8600 Software Release 7.2.0.0. Those limitations apply to this software release as well, unless noted otherwise.



Documentation Corrections

None

Copyright © 2013 Avaya Inc - All Rights Reserved. The Ethernet Routing Switch 8100/8300/8600/8800 is a trademark of Avaya, Inc.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Avaya.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Avaya Support on the web at <https://support.avaya.com/>
