



Ethernet Routing Switch 8800/8600

Readme

Software Release 7.2.11.0



Table of Content

Ethernet Routing Switch 8800/8600	1
Readme	1
Software Release 7.2.11.0.....	1
Table of Content.....	2
Software Release 7.2.11.0	3
Important Notices.....	3
Supported Hardware for Release 7.2.11.0	3
Upgrade Notes.....	3
Important Operational Notes.....	5
File Names for This Release	6
Version of Previous Release.....	7
Compatibility	8
SPB Multicast Interoperability	8
Changes in This Release	9
New Features in This Release.....	9
Old Features Removed From This Release.....	10
Problems Resolved in This Release	10
Known Limitations	13
Documentation Corrections.....	14



Software Release 7.2.11.0

Release Date: January 13, 2014

Purpose: Software maintenance release to address customer reported issues.

Important Notices

Supported Hardware for Release 7.2.11.0

Refer to the ERS 8800 Release Notes for a complete list of supported hardware - **Release Notes — Software Release 7.2.10.1 Avaya Ethernet Routing Switch 8800/8600 (NN46205-402, 11.07, Nov 2013)** available at <https://support.avaya.com>

Upgrade Notes

This section describes the Ethernet Routing Switch 8800 Software Release 7.2.11.0 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 7.2.10 Avaya Ethernet Routing Switch 8800/8600 (NN46205-400, Issue 09.02, November 2013)** available at <https://support.avaya.com>

Important Upgrade Note for systems using NLB configuration:

Due to a software issue present in releases older than 7.1.5.4 or 7.2.1.1 / 7.2.10.x, 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 addr> <mcast Mac addr> vid <vlan id> <ports>
exit
exit
exit
exit
exit
exit
exit
exit
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” operation is performed, the incorrect fields will get saved again and will need correcting.

This problem is not present in this release. However, if upgrading a system from a release older than 7.1.5.4 or 7.2.1.1/7.2.10.x directly to 7.2.11.0, 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) if upgrading from releases older than 7.1.5.2 or 7.2.1.1:

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 7.2.11.0 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 7.2.11.0 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 7.2.11.0 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.
- For 8812 ports utilizing LRM transceivers, link establishment may require an additional 15-20 seconds due to the synchronization/training time required for hardware linear mode equalization. The delays will be proportional to the number of LRM transceivers populating the lane.
- 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.
- Migration from PIM to SPB – When migrating a PIM based network to Multicast Service over SPB, one must wait 5 seconds between disabling PIM and enabling Multicast over SPB. Disabling of PIM results in removal of forwarding entries associated with PIM. This can take up to 5 seconds (timer). Not waiting for that interval to enable MCoSPB can result in the SPB-installed forwarding records to get deleted when PIM cleanup task runs. [wi01133076]
- Diagnostic features such as PCAP, Port Mirroring, DDM should be used cautiously as incorrect usage or over-usage of these features can adversely affect the device operations.
- The dos-check should only be run in a maintenance window.
- Ensure PCMCIA and flash have sufficient space available on them. Having large number of files on these or not having enough available space can cause system to malfunction.
- Live removal of PCMCIA / Compact flash is not supported and it can result in system reset. The “dos-stop /pcmcia” or “pcmcia-stop” command should be used to stop the device before it can be removed. Additionally before issuing the command, all user-driven and system-driven read/write activity to the device must be disabled by the user. Logging to PCMCIA, CLI Logging, SNMP Logging must be disabled and any manual or automated commands that would result in a write to PCMCIA should be blocked before running the “dos-stop /pcmcia” or the “pcmcia-stop” command.



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_72110.tar.gz	69,342,030
Ethernet Routing Switch images			
Boot monitor image for 8692 SF/CPU	8692 CPU and switch fabric firmware	p80b72110.img	1,198,466
Boot monitor image for 8895 SF/CPU	8895 CPU and switch fabric firmware	p80be72110.img	1,268,103
Run-time image for 8692 SF/CPU	Run-time image for 8692 SF/CPU	p80a72110.img	17,023,716
Run-time image for 8895 SF/CPU	Run-time image for 8895 SF/CPU	p80ae72110.img	15,957,770
Run-time image for R modules	Run-time image for R modules	p80j72110.dld	1,913,360
Run-time image for RS modules	Run-time image for RS modules	p80k72110.dld	1,982,108
Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz)	Image for the SuperMezz Card	p80m72110.img	17,105,210
3DES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c72110.des	56,124
3DES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce72110.des	51,972
AES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c72110.aes (this image includes the 3DES image)	27,644
AES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce72110.aes (this image includes the 3DES image)	25,316
MIB	MIB files	p80a72110.mib	5,681,229



MIB (zip file)	Zip file containing MIBs	p80a72110.mib.zip	897,042
MD5 checksum File	md5 checksums of all Release 7.2.11.0 software files	p80a72110.md5	1,249
FOQ for R modules	Feedback output queuing FPGA firmware	foq267.xsvf	5,320,469
BMC for R modules	BAP memory controller FPGA firmware	bmc776.xsvf	2,640,266
DPC for R modules	Dual port Controller FPGA firmware	dpc194.xsvf	2,642,001
PIM8630GBR	Programmable I/O module FPGA firmware; for the 8630GBR only	PI_769.xsvf	2,284,578
Firmware for RS modules	Contains FOQ, BMC, DPC, mirroring, and loopback images	rs_dpm_fpga.bin	4,538,368
PIM images for RS modules	PIM FPGA firmware required for 8612XLRS module only	pim8612XLRS.bin	60,183
	PIM FPGA firmware required for 8634XGRS module only	pim8634XGRS.bin	78,173
	PIM FPGA firmware required for 8648GBRS module only	pim8648GBRS.bin	79,891
	PIM FPGA firmware required for 8648GTRS module only	pim8648GTRS.bin	54,441
Trace files			
MPLS trace file	Trace file for MPLS. This is auto generated and appears on the PCMCIA after upgrade.	nbpdtrc.lo0	variable
EDM Help files			
EDM help files	Help files for EDM GUI	ers8000v7210_HEL P_EDM_gzip.zip	3,995,737
ERS 8000 EDM plug-in for COM			
EDM plug-in for COM	EDM plug-in for COM	ers8000v7.2.11.0.zip	5,336,310
Copyright	File containing copyright information	ERS8k.7.2.11.0_Cop yright.pdf	266,541

Version of Previous Release

Software Version **7.2.10.1**



Compatibility

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

Note: Updated EDM Help files and a new EDM plug-in version (v7.2.11.0) for COM are included with this release. Those must also be updated as part of upgrading to release 7.2.11.0.

SPB Multicast Interoperability

Multicast BEB Interoperability

- ▶ ERS8800 release 7.2.11.0 as a Multicast BEB interoperates with the following
 - Multicast BEB: ERS8800 release 7.2.0.0 or later
 - Multicast BEB: VSP9000 release 3.4.0.0 or later
 - Multicast BEB: VSP4000 release 3.1.0.0 or later
 - BCB: All except ERS8800 release 7.1.x
- ▶ If the Ingress BEB is a ERS8800
 - Minimum TTL=1 for UNI → NNI IP Multicast traffic
 - UNI → NNI Traffic will NOT be dropped if the TTL = 1
- ▶ Please note that the Ingress Multicast BEB will not modify the TTL on any product or release.

General TTL Handling Behavior

- ▶ For UNI → UNI Traffic if the TTL is 1
 - If the ingress and egress vlans are the same traffic will be allowed.
 - If the ingress and egress vlans are not the same traffic will be dropped.
- ▶ For NNI → UNI Traffic if the TTL is 1
 - Traffic will be allowed if the VSN has "igmp snooping" enabled
 - Traffic will be dropped if the VSN has "spb ip-multicast" enabled
- ▶ TTL will not be decremented for
 - UNI → NNI traffic
 - NNI → UNI traffic if the L2VSN has "igmp snooping" enabled
 - UNI → UNI traffic if the ingress and egress vlans are the same
- ▶ TTL will be decremented for
 - NNI → UNI traffic if the VSN (includes GRT) has "spb ip-multicast" enabled.
 - UNI → UNI traffic if the ingress and egress vlans are not the same

ERS 8800 Behavior Changes:

Multicast BEB L2VSN IP Multicast Traffic : TTL=1

- ▶ Traffic is dropped on the Egress BEB if all of the conditions below are met
 - A L2VSN has IGMP snooping enabled
 - Sender is connected to the SPB Network via a BEB that is either ERS8800 7.2.11.0 or VSP9000 3.4.1.0
 - Receiver is connected to the SPB Network via a BEB that is either ERS8800 prior to 7.2.11.0 or VSP9000 prior to 3.4.1.0
 - Traffic arrives at SPB Network ingress with TTL=1
- ▶ Such traffic would not be dropped for any other combinations of ERS8800/VSP9000 releases.

Multicast BEB ERS8800/VSP9000: Routed Multicast : TTL=1

- ▶ Traffic is dropped on the Egress BEB if all of the conditions below are met
 - Routed multicast (spb ip-multicast) is enabled on a L2VSN
 - Either the Sender or the Receiver is connected to the SPB Network via ERS8800 7.2.11.0
 - Receiver is on the same L2VSN as the Sender



- Traffic arrives at SPB Network ingress with TTL=1
- ▶ In previous releases traffic would not be dropped for such a combination of (Sender, Receiver, TTL)

Multicast BEB ERS8800/VSP9000: Routed Multicast TTL Decrement

- ▶ TTL is decremented on the Egress BEB if all of the conditions below are met
 - Routed multicast (spb ip-multicast) is enabled on a L2VSN
 - Either the Sender or the Receiver is connected to the SPB Network via ERS8800 7.2.11.0
 - Receiver is on the same L2VSN as the Sender
- ▶ In previous releases TTL would not be decremented for such a combination of (Sender, Receiver)

Changes in This Release

New Features in This Release

- RED KHI alarm for DRAM utilization has been introduced. When the DRAM utilization goes to 90%, an alarm "Health Status for DRAM USAGE is Red" is displayed on the console and also logged. User attention is required if DRAM usage reaches to the RED level. . If the condition persists, a system reset must be performed to recover the memory back to normal levels to prevent system malfunction. Note that the YELLOW alarm, which has been implemented in earlier release, is raised when DRAM utilization reaches 75%. [wi01136076]
- The command "show khi info" has been added to "show tech" output to enable capturing of KHIs as part of standard data collection. [wi01138265]
- Added support to allow the selection of CX SFP+ based links as part of MLT configuration [wi01113407]
- Added support to SONMP to recognize new VSP9010 chassis type [wi01120693]
- Support for Directed Broadcast to L2VSN has been added. [wi01115622,wi01119551]
- Added a check to reject values in the range 0 through 0x5DC for VLACP Ether Type parameter. These are not valid values and could cause issues if configured accidentally. This would not have been a valid operational configuration anyhow therefore no backward compatibility is provided [wi01124284, wi01105581]

Corresponding error message in EDM has also been corrected to state the valid range of VLACP Ether Type as 1500 to 65535 (0x05dc to 0xffff) instead of 0 to 65535. [wi01128200]
- To provide visibility into CP-bound packet storm scenarios, an automatic mechanism has been added to collect the CPP packet statistics data when the CP-bound path approaches congestion levels. This data will get written to the log file. The logging will automatically stop when the congestion subsides. [wi01125118]
- VLAN name is displayed truncated in ACLI commands such as "show vlan basic", "show vlan advanced" or CLI commands "show vlan info basic" and "show vlan info advanced". To overcome this limitation, a new command "show vlan name" has now been introduced to display the complete name of the VLANs in both CLI and ACLI mode. [wi01059893]
- Ability to enable IP-Multinetting under the VLAN Insertion screen was missing in EDM. The support for the corresponding MIBs (rcVlanIpMultinettingEnable and rcVlanAssociatedIpMultinettingVlan) has been added and the ability has also been added to the VLAN insertion screen. [wi01127837, wi01108367]
- Ability to turn on power for a specific slot using CLI/ACLI commands has been added. [wi01120228, wi01089953]
- Ability to configure RIP send and receive settings ("ip rip receive" and "ip rip send" commands) using the port-list option ("interface gigabitEthernet <port-list>") has been added. [wi01125610]
- An easily understandable error message "Root directory is full" was added to indicate when the flash was full. [wi01133079]



- ACLI OSPF redistribute and accept policy command syntax has been changed. Instead of the keyword “route-policy” the keyword “route-map” has been added to be consistent with other protocols. Older syntax will be automatically accepted for backward compatibility. [wi01100288]
- Ability to configure Soft Reconfiguration IN / OUT has been added to the BGP Peers tab in EDM.
- DFO changes have been added to capture the correct core dump in case of HA memory full scenario. This can aid in the analysis of field-found HA memory related issues. [wi01116495]

Old Features Removed From This Release

None

Problems Resolved In This Release

- Running certain large-output show commands with the more flag set to false can cause BFD sessions to flap. [wi01078538, wi01055558]
- Handling of PCMICA insertion/removal operation in the case of 8692 with supermezz CPs has been improved to avoid starving of the CPU for extended periods of time. [wi01102233]
- SLPP BPDUs were getting processed even when a port was in Alternate Discarding state. This has been corrected to instead drop the SLPP packets in the data path when the ports are in Alternate Discarding state. [wi01095952]
- The syntax for route-map configuration commands in ACLI mode had changed between 5.1.x and 7.x releases. During an upgrade, this could lead to configuration load errors and loss of the route-map configuration as well as configuration commands that followed the route-map section in the config file. To avoid this problem, a provision has been made in this release to recognize and read in the older versions of route-map syntax and convert it to the latest format automatically. This issue is not encountered for CLI configuration mode [wi01119893, wi01114905]
- In certain circumstances when a loop is created in a network, VRRP packets that are also being looped could overwhelm the receiving and processing of other control traffic including SLPP packets which are used to discover the loop, thus resulting in prolonging the discovery and resolution of the loop conditions. [wi01125993, wi01055419]
- When RSMLT is disabled on a node but left enabled on the IST-peer node, the static default route entry on that node was getting deleted from the hardware records. [wi01123837]
- Support has been added for the “-s” option for the ACLI mode command ‘show ip arp -s x.x.x.x y.y.y yrf xxx’ for the non-global-router case. The global router case already supported this option. Similarly the CLI mode command also already supported this option. [wi01125447]
- In ACLI mode the commands “ip bgp restart-bgp neighbor x.x.x.x vrf xyz” and “ip bgp restart-bgp vrf <vrf name>” did not restart the BGP neighbor instance under the particular VRF.[wi01124353]
- Improvements have been made to the internal logic pertaining to I2C accesses on the system bus. This will reduce the likelihood of I2C bus hang/reset conditions that earlier may have resulted in I2C access from multiple requesting processes. [wi01093101]
- When IPv6 traffic is transported via IPv4 tunnels, in certain instances packets could be forwarded to the wrong subnet due to wrong update of the encapsulation tunnel record. [wi01128126, wi01120076]
- With TACACS+ authentication, password length for FTP is now limited to 127 characters. If passwords greater than 127 characters are supplied, an error message is logged to the console and the log file and a failure status is returned to the FTP client. [wi01132511]



- SNMP MIB walk for the MIB 1.3.6.1.3.61.1.1.4.1 caused an “OID not increasing” error and would result in a premature exit from the MIB walk. [wi01132643]
- The SNMP MIB walk for RouteTable would fail with “OID not increasing” error when there was an alternate route present in the route table. [wi01120477]
- The GET-BULK operation for RcVrrpTmpRouterStatisticsTable was failing since ‘rcVrrpExtOperationsGloballpAddr’ is no longer a valid attribute name. ‘rcVrrpExtOperationsGloballpAddr’ has now been marked as obsolete allowing the GET-BULK operation to succeed. [wi01121424]
- In certain cases, execution of the L2TRACEROUTE command could cause line card reset on the intermediate hops on the traced path. [wi01119551]
- SSH and ACLI related robustness improvements have been made to address session lockup conditions mostly seen with continuous ssh login and logout activity. [wi01128500, WI01110085]
- There is a small window after a port is removed from a VLAN, when sender records might be forwarded to CP for processing. This is observed both in Multicast over L2VSN and L3VSN. A check has been added to ensure that new sender record will be created if and only if the port from where it is received is an active VLAN member. [wi01130217]
- TFTP related robustness improvements have been made to the allocated stack size in order to avoid potential for a crash due to stack overflow [wi01124288, wi01062093]
- ERS was not able to authenticate OSPF packets with LLS block since it was not fetching the correct md5 key from the incoming packet. This fix ensures the MD5 key is fetched from the correct location from the incoming packet, ignoring the LLS block. Note that LLS options is not supported on ERS. [wi01124890, wi01104497]
- When the route 0.0.0.0/x and x>0 is advertised to ERS node and withdrawn, it deletes the static default route entry from hardware. This fix ensures 0.0.0.0/x is not added to route table, and hence doesn't corrupt valid default route. [wi01126773]
- Receiving a LSA update packet with age greater than 3600 sec could lead to OSPF timer list corruption. A robustness fix has been added to drop the LSA packet with age > 3600 sec. [wi01129507]
- A robustness fix has been made to a SSH logout scenario where previously a CP crash could have resulted due to a NULL pointer access. [WI01123818, wi01127233]
- While extremely rare, it has been observed that a single egress queue may become locked in a way such that it will no longer forward data. A detection and recovery mechanism has been implemented to address this issue. Upon detection, a log entry “LCD F2E phase 2 reset lane #. Queue stalled (code 7)” will be displayed and the recovery routine will be invoked to clear the issue. This condition may be falsely reported (per lane) during cold boot in the presence of traffic and may be safely ignored. [wi01133211, wi01111582]
- Intermittent port flaps were observed on the 8812XL card due to a bad return value from a PHY's INTSTAT register read. [wi01122067]
- A crash can result in the SPB module, when an NNI interface on the device is frequently bouncing. [wi01121428]
- When a large number of IGMP joins and leaves occur, system memory could be depleted due to a memory leak scenario. [wi01135472]
- On configuring SPB NNI interface on 10G IO card, in some cases, STG 63 showed its root port as the 10G port rather than pointing to CPP. STG 63 is a special internal STG used in SPB scenario. [wi01121048]
- The check to block invalid mac from being set as IS-IS system-id was not working for the case where SPBM was globally disabled. The check has now been extended to take care of that case as well. [wi01128587]
- If a disabled route-policy were to be assigned to the IS-IS redistribution, then all IS-IS IP routes redistributed would have their external metric reset to 1 instead of using the metric from the routing table they were redistributed from. [wi01087399]



- A check added in 7.2.10.0 to prevent the use of invalid IP addresses in the “destIP” ACL command incorrectly also blocked valid Multicast IP addresses from being used in the “destIP” ACLs. This issue has been corrected in 7.2.11.0 as well as 7.2.10.1. [wi01121249, wi01121461]
- In HA mode, OSPF accept policy for VRF was not synced to standby CPU correctly. This would cause problems with saving the configuration on the new master after failover. The VRF ID being passed was incorrect. [wi01137535, wi01137394]
- SSH or console lockup could result if large quantity of text was copied and pasted into an active ssh session. A robustness fix has been added to prevent such lockup. [wi01135817, wi01134547]
- Executing “show ipv6 interface icmpstats gigabitEthernet” command could result in the session getting locked up. Change has been made to get the stats only for non-virtual interfaces thereby preventing such lockup. [wi01127765, wi01124332]
- ARP entries for next hop I-SID vlans were pointing incorrectly to a down interface instead of pointing to the I-SID. [wi01137527]
- A robustness fix to prevent failure or crash resulting from SNMP walk of IS-IS related MIBS has been added. [wi01121078]
- The value for Nickname displayed as output for CLI commands “show isis spbm i-sid” and “show isis spbm nick-name” is not consistent with the values displayed in EDM. [wi01139940]
- When two SPB systems are configured with nick-names which have the same first 4 bits, the code cannot detect the difference between the two nick-names because only the last 20 bits were used for the compare operation. [wi01139935]
- An SPB IPVPN default route can be lost from the hardware records of already operational line cards when a new line card is inserted into that chassis or is brought up. The situation can be recovered by disable/enable of ipvpn for that VRF. [wi01142304]
- Far end of a link to an 8683 10G port can take from 2 to 7 minutes to go to an “UP” state. The 8683 port disable / enable action included a lane reset that could leave the port in an hold down state. This reset is not required and has been removed. [wi01142817]
- Auto-negotiation configuration was being incorrectly applied to 100FX ports causing the 100FX ports to not re-establish link after a link bounce or link failure. As 100FX ports are not auto negotiated but operate at a fixed speed of 100M, auto-negotiation should not be applied to them. [wi01142536]
- HA crash could result from execution of an internal function meant to sync port statistics between the master and standby CPs. However this syncing is not required for PR cards as the COP maintains the statistics which are polled by CP. This redundant function call has now been removed to avoid the crash. [wi01136955]
- If an operationally down SPB NNI port is removed from an MLT, the port config is removed and port gets set to default configuration, when it should not. [wi01144241]
- A crash can result when deleting an OSPF interface with the command “config ip ospf interface delete”. [wi01143982]
- Output for the “show ip vrrp address vrf <vrfname>” command incorrectly displays the total number of VRRP Address as zero instead of showing the actual total number (e.g. “5 out of 0 Total Num of VRRP Address Entries displayed”). [wi01137079]
- L3VSN traffic with TTL=1 does not get dropped on the egress BEB. Instead TTL is decremented and TTL=0 traffic gets forwarded. This issue has been fixed. TTL=1 traffic will now correctly get dropped on the egress BEB. [wi01146161]
- QoS / ACL filter does not remark SPB BVLAN PBITS even when remark-dot1p is configured. [wi01128533]
- IGMP fast leave related configurations may not be loaded correctly during the config load. Setting fast-leave on a vlan interface does not always take effect, and may remain incorrectly as disabled. [wi01124300]



- EDM unknown prompt response “410 Not Authorized” was being displayed while logging out of the session in Internet Explorer. [wi01128197]
- Upon attempted logout from EDM using the Logout button, one is prompted whether or not to save the configuration file. If this action is cancelled, a dialog box popup “Not Authorized” is displayed. [wi01085722]
- In EDM, port QOS egress queue stats (“Egress Queue Set” tab) were shown under the wrong hierarchy - Edit-Port-IP. These have now been moved to the Graph-Port hierarchy. [wi01130478]
- Help message under “spanning-tree port <port> stp <stg> learning” was displayed incorrectly. This has been corrected to say “This command is Deprecated/Obsolete. Use 'spanning-tree stp <stp-id> learning fast'”. [wi01120224]
- If the default black-hole route is the only active route in the routing table, running “show ip route” command can result in the command output looping with repeated displays of that route. [wi01113152]
- A crash can result when “show ip ospf lsdb” is executed while the lsdb is being updated. [wi01123650]
- Traceroute command does not work for local interfaces. [wi01080999]
- In ACLI mode, when “no ip alternative-route” command is issued for the Global Router VRF, the configuration also gets applied to all VRFs. It should only be applied to the VRF in question. [wi01116502]
- Following an HA fail-over of a system, a few multicast streams did not recover. [wi01117672]
- For a scenario where there are two SPB paths to an end BEB, if the NNI ports towards the secondary path are bounced, quickly followed by bouncing the primary path, the correct port is not updated in the fast path, resulting in traffic loss after the primary path comes back up. [wi01117756]
- When the last port in a routed multicast over SPB VLAN is disabled the IP Multicast records may not get cleaned up correctly. These inconsistent records could prevent stream traffic from reaching newly joined receivers in that (S,G) which had that VLAN as an egress VLAN. [wi01120440]
- Bringing up 64th IS-IS adjacency would result in an error and not get added. [wi01124772].

Known Limitations

- 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
- 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.
- The SSH Max username length will accept only till 31 characters when RADIUS/TACACS enabled.
- The maximum username and Password length when TACACS is enabled with FTP authentication is 127 characters.
- Special character sequences in passwords can cause a crash. While configuring commands under “snmp-v3” like group-name, group-access, community-name, usm name and md5/priv passwords, it is advised NOT to use “%” symbol for names or passwords.[wi01137738]
- CLI crash is sometimes seen after executing the “show fulltech” command. Recommendation is to not use “show fulltech” and to use “show tech” and other show commands instead. [wi01136015]



- With terminal more false, “show fulltech” command will stop at “show ip bgp vpnv4” related commands(with entries on them present). Recommendation is to not use “show fulltech” and to use “show tech” and other show commands instead. [wi01135979]
- If an SSH session is closed while TACACS+ authentication is in progress, the session may not be fully terminated and still counts as one of the 8 allowed simultaneous SSH sessions. [WI01130627]
- A crash may occur when deleting a TACACS+ server while multiple TACACS+ authentications are in progress [WI01144526]
- It is possible to lose access to the console and other command line sessions when SSH and other applications are involved in rapid I/O operations. (WI01145067)The existing BGP Multi-Exit Discriminator (MED) value in BGP Announce Policies to eBGP Peers cannot be modified. The recommendation is to have IN Route policy on the neighbor Router to set the MED to any desired value, rather than having OUT route policy on the advertising Router. [wi01139882]
- A lower priority IPv6 VRRP hello causes backup to master transition. During this transition, the log message displays the cause incorrectly as “higherPriorityAdvertiseReceived”. This message can be ignored. It does not affect the functionality of VRRP6 protocol in the system.[wi01142757]
- IPv6 VRRP packet takes the NORMAL queue (RXQ# 4) and IPv4 VRRP packets takes the CRITICAL queue (RXQ# 1) to reach CPP. This can result in IPv6 VRRP transition for vrrp6 fast advertisement values during high CPU condition. [wi01142753]
- The command “ip bgp restart-bgp soft-reconfiguration <value>” is not available for GRT to initiate soft reconfig functionality in both CLI/ACLI. The recommendation is to configure “ip bgp restart-bgp neighbor x.x.x.x soft-reconfig <value>” to initiate the soft-reconfiguration per neighbor.
- The command “ip bgp restart-bgp vrf <vrf id> soft-reconfig <value>” expects IPVPN BGP related configurations to be enabled. This command initiates the soft-reconfiguration for all BGP instances present in that particular VRF. This is seen in both CLI/ACLI modes. The recommendation is to configure “ip bgp restart-bgp neighbor x.x.x.x vrf <vrf id> soft-reconfig <value>” to initiate soft-reconfiguration per neighbour.This command does not expect IPVPN to be configured. Note: The “ip bgp restart-bgp” command for GRT/VRF works fine without any difference in behaviour.
- The “conf ip arp static-mcastmac add” command needs to be entered twice for the static-mcastmac to show up in the table. Even though the configuration gets saved to the config file, after a reboot the static-mcastmac entry is not added to the table while the config is present in the config file. [wi01125450]
- A full SNMP walk on all MIBs fails with “OID Not Increasing” errors on certain MIBs. [wi01132644, wi01136185]
- On executing an SNMP walk, some MIBs timeout with the “Timeout: No Response from Server” error when the time taken by the switch to respond to the SNMP request is longer than the timer configured on the SNMP client. [wi01132637]

Please refer to the Known Limitations section of the Release Notes for Ethernet Routing Switch 8600 Software Release 7.2.10.0 and 7.2.10.1. No other known limitations have been found than those noted above.

Documentation Corrections

None



Copyright © 2014 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/>
