



Ethernet Routing Switch 8600

Readme

Software Release 5.1.7.0



Table of Contents

Ethernet Routing Switch 8600	1
Software Release 5.1.7.0.....	1
Table of Content.....	2
Software Release 5.1.7.0.....	4
Important Notices.....	4
Platforms Supported.....	5
Notes for Upgrade	5
File Names for This Release	6
Version of Previous Release	8
Compatibility	8
Changes in This Release	9
New Features in This Release.....	9
Old Features Removed From This Release	9
Problems Resolved in This Release.....	9
Outstanding Issues.....	11
Known Limitations	15
Documentation Corrections.....	11
Software Release 5.1.6.0.....	12
Changes in This Release	12
New Features in This Release.....	12
Old Features Removed From This Release	12
Problems Resolved in This Release.....	12
Documentation Corrections.....	16
Software Release 5.1.5.0.....	17
Changes in This Release	17
New Features in This Release.....	17
DfO KHI	17
ACLI.....	17
Old Features Removed From This Release	17
Problems Resolved in This Release.....	17
Software Release 5.1.4.0.....	18
Changes in This Release	18
New Features in This Release.....	18
Old Features Removed From This Release	20
Problems Resolved in This Release.....	20



Software Release 5.1.3.0	23
Changes in This Release	23
New Features in This Release.....	23
Old Features Removed From This Release	23
Problems Resolved in This Release	23
Software Release 5.1.2.0	26
Changes in This Release	26
New Features in This Release.....	26
Old Features Removed From This Release	26
Problems Resolved in This Release	26
Software Release 5.1.1.1	30
Summary	30
Changes in This Release.....	30
Software Release 5.1.1.0	32
Changes in This Release.....	32



Software Release 5.1.7.0

Release Date: November 18, 2011

Purpose: Software maintenance release to address externally found customer issues.

Important Notices

Avaya Inc has acquired the Enterprise Solutions business from Nortel. This acquisition includes the ERS 8600 and software described in this document. During the transition of all assets and support infrastructure of products and services to Avaya, Nortel provides certain infrastructure support. As such, you will find the Nortel name, and pointers to Nortel support locations on the Internet being referenced intermixed with Avaya in this document.

REGARDLESS OF SOFTWARE VERSION, the system-monitor flag should be checked on all systems to be sure it is enabled. This flag should always be enabled as it enables a software monitoring capability to detect and respond to abnormal software loop conditions. The flag setting can be checked via the command “show config”. NOTE: Enabling this flag only takes effect after a reboot and must be saved in the config prior to reboot.

The display should be similar to the following:

```
ERS8600:6# show config
Preparing to Display Configuration...
```

```
#!flags m-mode false
#!flags enhanced-operational-mode false
#!flags vlan-optimization-mode false
#!flags global-filter-ordering false
#!flags r-mode true
#!resource-reservation max-vlan false
#!resource-reservation multicast 2048
#!flags multicast-check-packet true
#!flags system-monitor true (enabled)           or potentially false (disabled)
#!flags regular-Autoneg false
#!record-reservation filter 4096
```

If the system-monitor flag is set false, then it should be changed to true which can be accomplished through JDM by following the menu option sequence “Edit Chassis” -> “System Flags” and then look under “System Monitoring” at the bottom of the screen. Checked equals enabled.

To set via SNMP use:

```
MIB is rapidCity.rcMgmt.rcChassis.rcChasSystemMonitorEnable
```

```
snmpset -v 1 -c public x.x.x.x enterprises.2272.1.4.41.0 1
```

Where x.x.x.x = some IP address associated with the switch.

To view the setting via SNMP use:

```
snmpget -v 1 -c public x.x.x.x enterprises.2272.1.4.41.0
```



Output is either:

FALSE - SNMPv2-SMI::enterprises.2272.1.4.41.0 = INTEGER: 2 (disabled)

TRUE - SNMPv2-SMI::enterprises.2272.1.4.41.0 = INTEGER: 1 (enabled)

Platforms Supported

All Ethernet Routing Switch 8600 modules are supported in the 8006, 8010, and 8010co Chassis. Full slot support for all modules may be dependent on the presence of the High Performance (HP) Backplane. There is an orderable (and chargeable) upgrade option for the HP Backplane.

The following modules are not supported in the 8003 chassis:

- 8692SF/CPU
- All R/RS modules

NOTE: R/RS-series modules are supported in the 8010co chassis only with a High Performance Backplane installed.

Please refer to the following documents for extra details and as reference material in addition to this Readme:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 5.1.0.0 (Doc # NN46205-402, Rev 3.01)
- Readme for Ethernet Routing Switch 8600 Series Switch Software Release 5.0.5.0. Release 5.1.2.0 has equivalent CR fix functionality as Release 5.0.5.0.
- Nortel Ethernet Routing Switch 8600 5.1 Upgrade Manual (Doc # NN46205-400, Rev 3.01)
- Nortel Ethernet Routing Switch 8600 Installation – Modules, for Software Release 5.1 (Doc # NN46205-304, Rev 3.02)
- Nortel Ethernet Routing Switch 8600 Administration, for Software Release 5.1 (Doc # NN46205-605, Rev 2.03)
- Nortel Ethernet Routing Switch 8600 Routine Maintenance, for Software Release 5.1 (Doc # NN46205-312, Rev 2.01)

Notes for Upgrade

Please see Upgrade Guide and Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 5.1.0.0 available at www.avaya.com/support.

NOTE: If upgrading to 5.1.7.0 code within any SMLT designed network from a release prior to 4.1.8.2 or 5.0.1.0, i.e. a release which runs the 'older' SMLT architecture, then care should be taken to follow the proper upgrade steps. Please review the 5.1.0.0 or the 5.0.1.0 Release Notes for specific details.

NOTE: If upgrading to 5.1.7.0 code from a release prior to 5.1.0.0 with the traffic filter configuration having DSCP value 0x28, it is required to modify this DSCP value manually to 0x2E again after the upgrade.

NOTE: dpc194.xsvf is a new DPC FPGA image for R-modules that can only be used with post 5.1.3GA release. Do not use this DPC FPGA image with pre5.1.3 releases; it may result in R-modules not booting up successfully and one may need RMA to get the card back in good state.

NOTE: While upgrading from an earlier release of 5.1.x to 5.1.6.0 or later, traffic filter configuration having DSCP value 0x28(e.g. 40) will be changed to 0x2E(e.g. 46) according to ieee8021p value configured. In the latest release of 5.1.x.0, the DSCP value of 40 is mapped to ieee8021p value of 5 (instead of 6 in older release). Hence, it is required to reconfigure the DSCP value manually back to 0x28(e.g. 40) post the upgrade.



File Names for This Release

Module or file Type	Description	File name	Size in bytes	
Software tar file	Tar file of all software	Deliverables (includes images that also contain encryption software)	pr86_5170.tar.gz	62618927
Ethernet Routing Switch images				
Boot monitor image	CPU and switch fabric firmware	p80b5170.img	1141319	
Run-time image	Run-time image	p80a5170.img	12802767	
Run-time image for R modules	Image for R modules	p80j5170.dld	1620988	
Run-time image for RS modules	Run-time image for RS modules	p80k5170.dld	1679884	
Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz)	Image for the SuperMezz card	p80m5170.img	12903482	
3DES	Encryption module for privacy protocol with Secure Shell (SSH)	p80c5170.img	55928	
AES	Encryption module for privacy protocol for SNMPv3. Includes AES and DES	p80c5170.aes (this image includes the DES image)	26947	
MIB	MIB files	p80a5170.mib	4391441	
MIB (zip file)	Zip file containing MIBs	p80a5170.mib.zip	694824	
MD5 checksum file	md5 checksums of all Release 5.1 software files	p80a5170.md5	1358	
Runtime image for ATM	Runtime image for the ATM module	p80t5170.dld	906024	
Runtime image for POS	Runtime image for the POS module	p80p5170.dld	701771	
Firmware images				
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
SSL images			
SSL cluster upgrade	Ethernet Routing Switch 8600 clustered SSL modules self-installing runtime image/upgrade	p80s5170.pkg	5988896
SSL boot monitor	Ethernet Routing Switch 8600 SSL module boot monitor	p80s5170.img	7528448
SSL upgrade instructions	Ethernet Routing Switch 8600 SSL upgrade instructions	p80s5170.upgrade	1481
SSL installation instructions	Ethernet Routing Switch 8600 SSL installation instructions	p80s5170.install	2895
SSL diagnostics	Ethernet Routing Switch 8600 SSL diagnostics	p80s5170.diag	19460381
WSM images for Ethernet Routing Switch 8600			
WebOS firmware image	WSM WebOS v10.0.34.0 firmware image	Wsm1003400_mp.img	845560
WebOS binary	WSM WebOS v10.0.34.0 binary image	Wsm1003400_bin.img	1376256
WebOS boot image	WSM WebOS v10.0.34.0 boot image	Wsm1003400_boot.img	43004
Device Manager images			
Solaris for SPARC image	Device Manager software image	jdm_6200_solaris_sparc.sh	
Microsoft Windows image	Device Manager software image	jdm_6200.exe	
Linux image	Device Manager software image	jdm_6200_linux.sh	
Service Delivery Module images			
Service Delivery Module Firewall	Boot ISO for the NSF Service Delivery Module Firewall booting from CD-ROM	NSF5100_2.3.7.0_SDM_R60.iso NSF5100_2.3.7.0_SDM_R65.iso	
	Upgrade package for the NSF Service Delivery Module Firewall	NSF5100_2.3.7.0_SDM_R60.pkg.gz NSF5100_2.3.7.0_SDM_R65.pkg.gz	
Service Delivery Module TPS	Boot image for TPS Intrusion Sensor	Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	
	Boot ISO for TPS Defense Center booting from CD-ROM	NortelTPSDefenseCenter_2x70v4.5.0_627_Install.iso	



	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1.	Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	
	IS upgrade download verification file.	Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh.md5	
Trace files			
MPLS trace file	Trace file for MPLS. This is auto generated and appears on the PCMCIA after upgrade.	nbpdtrc.io0	variable

Version of Previous Release

Software Versions **5.1.6, 5.1.5, 5.1.4, 5.1.3, 5.1.2, 5.1.1.1 and 5.0.5.0**

Compatibility

This software release is managed with Java Device Manager (JDM) release 6.2.1.3 or higher.

Note:

When a SNMP user needs only layer-1 access through JDM, including only layer-1 MIBS will not complete the mib-view access. Some minimum set of layer3 OID's are required in mib view to access the switch using JDM. This is the current design of JDM.

Below is the list of minimum set of OID's to access the switch using JDM. These OIDs need to be part of the mib-view in case we need to access JDM. Otherwise, JDM will give a throws java.lang.ClassCastException.

1.3.6.1.4.1.2272.1.203.1.1	rcVrfObjects
1.3.6.1.4.1.2272.1.4.8.1.1.2	rcChasPowerSupplyOperStatus
1.3.6.1.4.1.2272.1.100.6.1.2	rc2kCardFrontType
1.3.6.1.4.1.2272.1.100.1.1	rc2kChassisPortOperStatus
1.3.6.1.4.1.2272.1.100.6.1.5	rc2kCardFrontOperStatus
1.3.6.1.4.1.2272.1.4.10.1.1.2	rcPortType
1.3.6.1.4.1.2272.1.100.2.1.4	rc2kCpuEthernetPortOperStatus
1.3.6.1.4.1.2272.1.100.2.1	rc2kCpuEthernetPortEntry
1.3.6.1.4.1.2272.1.4.7.1.1.2	rcfanOperStatus

This software release supports the Web Switching Module (WSM) release WebOS 10.0.34.0. This code is found on Nortel web site under Content Networking -> Web Switches -> Ethernet Routing Switch Web Switch Module.

This software release supports SDM FW Release up to 2.3.7.0. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660, which provides a link to CheckPoint web site.

This software release supports SDM TPS Release up to 4.7.0.2. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.



To download any of these code releases requires valid Nortel support web access, as well as valid CheckPoint support site access for most FW code.

Changes in This Release

New Features in This Release

None.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Platform

- CPU lockup is possible on abnormal termination of SSH sessions with Cli logging enabled. (wi00869014)
- When a MAC address frequently moves between ports or MLT's, the associated MAC address entries in the FDB and the ARP table may not always be in sync. (wi00885468)
- Configuration change to "DDM Port Down" might bounce all the active ports. (wi00890497, wi00885004)
- Line cards 8648 GBRS, and 8634 XGRS do not send Remote Line Fault (RLF) clear signal. (wi00875414, wi00898305)
- The CPU might lockup with continuous Ring Buffer error messages on the console. (wi00927025)
- The cause of 10Gig port's link state transition is not captured in the log. (wi00884896)
- CurrentMacEntries counter, used for limiting FDB learning, may not have the actual count of learned MAC entries. (wi00908755)
- ICMP Replies sent on behalf of the RSMLT Peer are not always sent with the peers IP/MAC address. (wi00936870)
- A failover can occur on receipt of an inter-process message when attempting to free a part of a data structure after the structure itself is freed. (wi00876963)
- Routed NLB multicast traffic will not egress an R/RS-module VLAN tagged port if that port belongs to both the ingress and egress VLAN. (wi00888267)
- DDI functionality of the Supported Pluggables does not always working properly. (wi00891162)
- Hidden files such as snmp_usm.txt cannot be copied between master and backup CPU's. (wi00890818)
- DDI functionality for all Avago/Agilent Pluggables has been removed. (wi00930216)
- When a static SSM channel is deleted, the event is not synced to the slave CPU. Hence, after an HA fail-over, the static SSM channel is present. (wi00869500)
- On a triangle SMLT setup with LACP enabled, an uninitialized LACP value can cause a core. (wi00952662)
- SMLT and IST HELLO messages were sometimes logged as SMLT_MSG_TYPE_INVALID (wi00875027)
- When VLACP is disabled, non IEEE reserved mac addresses are no longer bridged. (wi00949517)

- Spanning Tree path cost is not updated after a speed change on a copper port resulting from enabling auto-negotiation on the remote end. (wi00939938)
- The XFP type indicated in the log may occasionally be incorrect or incomplete after the part is removed and reinserted quickly. (wi00869323) (wi00908211)
- If LACP ports are configured in a STG other than STG 1, and an R module or RS module where on which LACP ports are located is reseated or reset either via JDM or CLI, the 8600 starts sending the LACPDU's tagged with VLAN 4095 which are then dropped by the other end, causing the LACP trunk go down. For default STG 1, this is not an issue. (wi00896461)

CLI/NNCLI

- Errors occur if "Max-Queue-Length" parameter is modified for the queues Q0, and Q7 of custom queue set Q8, and for the queues Q0, and Q63 of custom queue set Q64. (wi00874094, wi00906392)
- Disabling/Enabling BFD on BGP peer does not work when using the BGP Peer Group name. (wi00880674)
- When queue length of balanced queue of saved queue set is modified, active queue set disappears. (wi00874096)
- "Show" and "Config info" commands for DDM monitor display different state information for DDM. (wi00881819)
- For every new MLT entry created through JDM, a LACP entry gets created. (wi00898543)
- Using ACLI, an inter-vrf BGP to BGP redistribution can be configured. (wi00893869)
- The "show khi port" command only shows SMLT Port Down events, not SMLT Port Up events. (wi00844647)
- Using JDM, the highest priority queue can be configured with max-length less than the default max-length. (wi00906392)

QOS

- CPU generated Control packets (ICMP, DHCP, TELNET, SNMP, ARP, SSH, IPFIX) are marked as network critical traffic (QOS 7). (wi00870809) (wi00932051) (wi00939810)

IP Routing

- Reception and withdrawal of routes with a prefix 0.0.0.0 and mask not equal to 0.0.0.0 might corrupt the default route. (wi00931026)
- BGP peer connections may be lost when a 8608GBE I/O card is inserted or reseated. (wi00943009)
- If an OSPF External LSA with a non-zero forwarding address is received, routes from other External LSA's from the same source which have no forwarding addresses may be placed in the route table with that forwarding address. (wi00906719)
- Inter-vrf redistributed routes are not correctly synced to the slave CP. (wi00893873)
- Disabling OSPF route redistribution also disables RIP route redistribution (wi00893871)
- When a BGP peer group is created, that event is not synced to the slave CPU; this error is displayed: "BGP ERROR: rclpBgpPeerGrpTblSetBody:EventSync failed". (wi00938857)

- BGP Peer Group BFD Disable does not bring down the BFD session. (wi00904711)
- Static BFD sessions do not come up after a failover. (wi00859407)

SNMP

- SNMP traps are not sent by the Circuit less IP address though configured as sender-ip_ They are sent by interface IP address. (wi00894073)
- MIBs in ipMRouteEntry, ipMRouteNextHopEntry, ipMRouteInterfaceEntry do not return correct values. (wi00858085)
- In ACLI mode MIB walk from SNMP application might fail in some scenarios. (wi00921309)

MultiCast

- Enabling "more-ip-specific-non-local" command on RP PIM router might stop new multicast flows to the receivers. (wi00904417)
- On a PIMSM configuration with RP not in the SPT, initially a prune is sent towards the RP. If the topology changes so RP is in the SPT, a join is sent but is not propagated to and traffic does not recover. (wi00894381)
- A failover may occur when a multicast route is deleted if it is associated with an MLT port and no port is active in the MLT. (wi00940394)
- With two PIM RP routers running MSDP directly connected to a multicast source, both originate MSDP SA's even though one is recognized as the PIM DR. (wi00872502) (wi00733867)

Outstanding Issues

- In some MSDP configurations, an excessive number of MSDP INFO messages .may be logged. (wi00928703)
- Failovers may occur due to exceptions following TACAS authentication failures. (wi00945178)
- RFI is not working on 10Gig ports of 8634XGRS line cards. (wi00936425)
- Frequent COP SW Packet memory refresh errors may be logged. (wi00921495, wi00941598)

Documentation Corrections



Software Release 5.1.6.0

Release Date: March 04th, 2011

Purpose: Software maintenance release to address externally found customer issues.

Changes in This Release

New Features in This Release

None.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Platform

- Enabling the WLAN mode option on the 8683XZR card now functions correctly. While the enable/disable commands for WAN mode are displayed for the 8612XLRS and the 8683XLR card types, the system responds with appropriate error message indicating that WAN mode is cannot be enabled on these cards types. (wi00870205)



- Broadcast and multicast traffic that should be flooded within a VLAN is no longer forwarded out the same MLT as it arrived if the ingress port is in the port range 41-48 on an 8648GTR or 8648GTRS card. (wi00867614)
- Packets with an unknown SRC MAC were dropped in RSP in some packet learning scenarios where an IST link was configured on the first port of a lane. (wi00866355)
- HAL INFO GBIC inserted in slot x Port x Type:Gbic1310(Lx) message may appear in the log randomly when a 8630R module is installed. (wi00858345)
- Ports 41 to 48 of 8648GTR card would egress all the traffic only on Queue 0 regardless of DSCP marking or any port DiffServ settings. (wi00856177)
- A CPU reset may occur if the switch attempts to send TCP packets whose destination matches a static blackhole entry in the routing table. The static route is configured with a nexthop = 255.255.255.255 and pref = 255. (wi00854617)
- CP may crash when it is continuously reading or writing the port register of legacy E-Module while the line card is being physically removed from the chassis. (wi00853206)
- TCP port 111 no longer accepts connections in order to avoid potential security concerns. (wi00852411)
- When enabled, the "untagged-frame-discard" feature on a port now properly discards untagged broadcast frames if VLACP is enabled on the port. (wi00846353)
- A message buffer handling interaction has been addressed that could lead to the ERS8600 very infrequently generating a core file or appearing to hang during a reboot and display the log message "Master CP failed before table synchronization." (wi00845719)
- When a XD SFP is present on a port that is a member of an MLT, the ERS8600 will now allow additional ports housing different XFP types to be added to the MLT. Previously if the GBIC type in the other MLT ports was anything other than type "Gbic1310Xd" the system would generate the message " Error: MLT ports different types" and prevent any port housing Non XD SFPs from being added to the MLT. (wi00843343)
- Switch over to a warm standby CPU can fail if a legacy card is present and ARP packets are received during the memory test portion of the CPU initialization sequence. (wi00834470)
- When performing an snmpwalk of the MIB rcPortVLacpPortState, the first operationally up (VLACP UP) port will always show **DOWN**, instead of its correct UP state. This is a monitoring only issue and is not service impacting. (wi00819545)

CLI/NNCLI

- When accessing the ERS8600 via SSH in ACLI mode, using the option "q" to abort the display of paged command output will no longer prevent subsequent output for any commands in that session. (wi00884376)
- A Config file with a space in the string parameter of the msg-control force-msg configuration will not load properly during a reboot. (w i00876846)
- ERS 8600 no longer reboots or freezes when an MD5 command is entered using an invalid filename. (wi00853274)
- In ACLI mode, configuration changes to the RIP attributes, listen and supply are now properly saved in the configuration file. (wi00733748)

Filters

- ERS8600 with Filters configured no longer generate the software Error message, "COP-SW-FILTER ERROR Slot x: ercdFilterGetStatsCtrBlk: Failed!! Ingress Filter Stats Counter Memory is FULL". This error message is changed to trace message. (wi00847628).
- When enabling port mirroring on an RS module port where a traffic filter(ACL) is configured which specifically contains a "port add" clause, high levels of packet discards are no longer encountered on the data path. This issue was not encountered with using inVlan ACLs. (wi00852577)
- When using debug mirroring through a filter on an RS card traffic impact and card resets no longer occur. This issue was previously encountered intermittently if filter statements such as "filter acl <X> ace <Y> debug count enable mirror enable" are used without a corresponding destination in the ACE record. (wi00854540)
- QOS policed configurations will now allow the deletion of a configured ACL after HA failover. (wi00854642)
- Port mirroring mode can now be set to rxFilter separately when mirroring legacy ports. This mode is not applicable to R or RS module ports. (wi00856197)

IP

- An OSPF network configured with non-unique OSPF router id may trigger the ERS8600 to reset if it receives a NET_LINKS packet from a remote router whose ID is the same as one directly connected to the ERS8600. (wi00858330)
- The ERS8600's IP Spoof-detect feature does not properly protect the VRRP IP address in the network. (wi00857095)
- In a stable, converged network, VRRP backup node need not start the Hold Down Timer to become Master. The VRRP Hold Down Timer is now only started on node recovery. (wi00856647)

DHCP

- The ARP entry of an end device was being removed from the ARP table upon reception of a DHCP renew response packet (ACK) from the DHCP Server. This can cause a brief outage as the ARP entry needs to be re-populated at this time. Since this is not a new IP address, just a renewal of the IP address lease the ARP entry now remains in place. (wi00866471 and wi00866465)

IP Multicast

MLT / SMLT

- Addition and deletion of static Mac entries on an SMLT in SMLT up state may result in stale software and hardware Mac entries. (wi00867171)
- Traffic prioritization (QOS) when using MLT's with an ID > 96 default to QOS of 0. (wi00859025)

- When the destination MAC address is learned (unicast case), Packets coming on the IST ports should not be forwarded to SMLT when SMLT remote is up and packets come from the same SMLT. (wi00852563)

BGP/VRF

- Configuring BGP route redistribution requires a valid advanced or premier license if fewer than 10 BGP peers are configured. (WI00871534)
- A BGP peer remains in an idle state after a link state transition if running over a link with BFD enabled and the BFD holdoff timer is configured. (wi00863051)

PIM

- When multicast traffic is sent to two directly connected switches A and B, where switch A is the PIM designated router and switch B is connected to a third switch attached to a PIM receiver, a prune sent from switch B to switch A is now accepted. Previously this prune message was not handled properly and resulted in duplicate traffic on Switch B, (wi00850153)
- In a square router configuration (non-SMLT) where a prune message is forwarded toward the Rendezvous Point via a different router than the one in IPMC Source Path Tree, the prune is now properly forwarded and duplicate traffic is not encountered. . (wi00871265)

RADIUS

- When using RADIUS authentication, a password set to more than 16 characters may be truncated to 16 characters when the RADIUS packet is generated.(wi00850719)

Known Limitations

The three tables **ipMRouteTable**, **ipMRouteNextHopTable** and **ipMRouteInterfaceTable** are not supported in the ERS8600 5.1.6.0 release. As the implementation to fully support these MIBS and collect data for them is not present in the code, access to these MIBS will be removed from the next release.

Alternative functionality:

The ERS8600 provides the proprietary MIBS **rcIpMRouteStatsTable**, **rcIpMRouteHwEgressVlansTable**, and **rcIpMRouteHwSourcesTable** which support the similar functionality.

When a SNMP user needs only layer-1 access through JDM, including only layer-1 MIBS will not complete the mib-view access. Some minimum set of layer3 OID's are required in mib view to access the switch using JDM.This is the current design of JDM.

Below is the list of minimum set of OID's to access the switch using JDM. These OIDs need to be part of the mib-view in case we need to access JDM. Otherwise, JDM will give a throws java.lang.ClassCastException.

1.3.6.1.4.1.2272.1.203.1.1 rcVrfObjects



1.3.6.1.4.1.2272.1.4.8.1.1.2	rcChasPowerSupplyOperStatus
1.3.6.1.4.1.2272.1.100.6.1.2	rc2kCardFrontType
1.3.6.1.4.1.2272.1.100.1.1	rc2kChassisPortOperStatus
1.3.6.1.4.1.2272.1.100.6.1.5	rc2kCardFrontOperStatus
1.3.6.1.4.1.2272.1.4.10.1.1.2	rcPortType
1.3.6.1.4.1.2272.1.100.2.1.4	rc2kCpuEthernetPortOperStatus
1.3.6.1.4.1.2272.1.100.2.1	rc2kCpuEthernetPortEntry
1.3.6.1.4.1.2272.1.4.7.1.1.2	rcfanOperStatus

Documentation Corrections

A new "built-in" filtering access control template (ACT), VPS DEFAULT ACT (4080) has been added and will be included in an update to the **Configuration — QoS and IP Filtering Guide**.

The following note will be added to the **Installation — SFP, XFP, GBIC, and OADM Hardware Components Guide**

Digital Diagnostic Monitoring (DDM) is not available non-Nortel/Avaya SFPs.



Software Release 5.1.5.0

Release Date: March 04th, 2011

Purpose: Software maintenance release to address externally found customer issues.

Changes in This Release

New Features in This Release

DfO KHI

ACLI

Old Features Removed From This Release

None.

Problems Resolved in This Release

Platform

- If all the MGIDS are exhausted in the system and the port which is part of the vlan is enabled/disabled, then the MGID Errors are seen even though there are no provisioning errors encountered while configuring. This issue is resolved by avoiding the check for free MGIDs while enabling/disabling the port. (wi00732528)
- QOS mapping configuration is not in sync between Master and Slave after HA Failover. This is due to incorrect table sync for the QOS Mapping Parameters. This issue resolved (Wi00698375)

VLACP

- The “time-out” scale for VLACP protocol to bring the port down was miscalculated and taking twice the time of configured value. This issue have been rectified. (Wi00728743).
- The periodic timer interval for sending VLACP Hellos was miscalculated and taking twice the time of configured value. This issue have been rectified. (Wi00728749)



Software Release 5.1.4.0

Release Date: December 03rd, 2010

Changes in This Release

New Features in This Release

VLACP HOLD Enhancement

During SMLT node failure scenarios, traffic loss may be observed in certain scaled SMLT configurations with hundreds of SLTs, hundreds of ports and tens of VLANs. The root cause for the traffic loss was that the ERS8600 ports would come up prematurely at the physical layer causing the remote end to start sending traffic toward the ERS8600 that just came up. On the ERS8600 that just rebooted, the communication between the line cards and the CP may take several seconds in such scaled configurations. This resulted in black-holing the traffic arriving on such ports which were physically up but all operational configuration was not yet performed on those ports by the CP. The VLACP SUBTYPE HOLD feature introduces a new VLACP PDU with a new subtype HOLD to help reduce traffic loss in such scenarios.

The goal of this new implementation is to "hold down" all VLACP enabled links for a specific period of time after a reboot. This prevents remote VLACP enabled devices that understand the new VLACP HOLD PDU from sending data to the ERS8600. This will ensure that all VLACP enabled ports on the ERS8600 have had sufficient time to come up with all operational configuration and are ready to receive and forward the ingress traffic.

ERS8600 switches with 5.1.4.0 release are capable of both sending and receiving VLACP HOLD PDUs. Future code revisions of the Baystack switch family will support receipt and processing of VLACP HOLD PDUs, but will not generate them. Please refer to the applicable product release notes for information regarding product specific software levels required for support of this VLACP enhancement. VLACP is an Avaya proprietary protocol and hence this enhancement is not applicable when connecting to switches from other vendors.

By default, the VLACP HOLD feature will be disabled. The feature is enabled by configuring a positive value for VLACP HOLD Time. The VLACP Hold Time value configured should be selected based on the specific recovery implementation requirements, size and recovery characteristics for your network implementation.

The following new CLI commands are introduced to support:

```
ERS-8610:6/config/vlaccp#?
```

```
Sub-Context:
```

```
Current Context:
```

```
disable
```

```
enable
```

```
info
```

```
hold-time <seconds>
```

```
ERS-8610:6/config/vlaccp# hold-time ?
```

```
value in seconds for hold time
```



Required parameters:

<seconds> = vlacp hold time {0..60}

Command syntax:

hold-time <seconds>

```
ERS-8610:6/config/vlacp# info
                               Vlacp           : enable
                               Vlacp hold-time : 20
```

```
ERS-8610:6/show/vlacp# info
```

```
=====
                               Vlacp Global Information
=====
                               SystemId: 00:18:b0:5c:a0:00
                               Vlacp           : enable
                               Vlacp hold-time : 20
```

SLPP Enhancement

- **Introduced New EtherType for SLPP PDU:**

From this release EtherType 0x8104 is replaced by 0x8102 as the default EtherType for SLPP PDUs. SLPP EtherType can still be configured through CLI. Old EtherType 0x8104 will be processed to handle the backward compatibility and upgrade scenarios.

If the received ether type is old default ether type (0x8104) or current default ether type (0x8102) or currently configured ethertype, then the packet is classified as SLPP PDU.



The Ethertype can be displayed using following existing CLI command.

CLI Command:

```
config slpp info
```

```
Sub-Context:
```

```
Current Context:
```

```
                add :  
etherType (hex) : 0x8102  
                operation : disabled  
                tx-interval : 500
```

- **Re-Arm per port SLPP PDU Receive Counter issue:**

Currently per-port SLPP PDU receive counter is never reset, resulting in shutting down links wrongly after months of running when the counter hit the pre-defined limit. This issue has been addressed in this release by resetting the counter if the switch not received expected number of SLPP packets on the port in a certain period of time. This timer is set to 24 hours.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Platform

- 8648GTR and 8648GTRS modules do not negotiate speed and duplex as expected once Customizable Auto-negotiation Advertisements (CANA) has been used on to limit link speed on a port and is subsequently disabled on that port. This issue is resolved (wi00518718)
- When PCAP is enabled and disabled, a control bit in fast path was not cleared correctly. This results in packets received on the port that had been monitored through PCAP being forwarded to the CP incorrectly and potentially cause high CPU utilization. This issue is resolved (wi00564156)
- In prior releases, a timing issue existed in single CPU systems where physical reseating of the CPU card would result in a COP-SW Exception event to be copied to the CPU during switch recovery. To avoid this scenario, the chassis should be power cycled rather than reseating the single CPU. (wi00508448)
- Port based shaper functionality may be disabled in hardware for R/RS modules after a port state change. This issue is resolved (wi00564813)
- SNMP get failed on rclpfixExporterStatsTable since the length field passed in the OID is not supported by the agent code. This issue is resolved (wi00564237).
- For tagged ports, the QoS queue mapping for Layer 3 packets received on R/RS modules has been corrected to align with E-modules. (wi00518702).
- In prior releases, a very rare scenario exists where the system monitor could potentially incorrectly detect the tTrapD task to be in an infinite loop and initiate recovery action. This issue is resolved (wi00508402).

- If 'ctrl-C' is entered at certain points in the boot up sequence the normal boot sequence is aborted and the user is entered into the shell access level of the switch. Sending 'ctrl-C' no longer results in user access to the shell level. (wi00675102).
- TFTP packets egressing out of line card port are marked with dscp 0x30, instead of default value. This issue is resolved (wi00564778).
- Executing the 'show sys pluggable-optical-modules info' command no longer results in repeated GBIC inserted messages being generated. (wi00564844).
- PIM mroute flapping could potentially be encountered in networks configured with short timer values for the multicast forward cache timeout. This was only encountered for multicast streams with low data rates in MLT or SMLT topologies and was dependent on which link in the S/MLT the data stream was active. Mroute flapping will no longer be encountered in this scenario. (wi00564358)

CLI/NNCLI

- Executing the 'show fulltech' command in NNCLI mode in some configurations would previously lead to unexpectedly high CPU utilization. Optimizations in the execution flow have now been made to minimize the impact of running the show 'fulltech'. (wi00564789).

IP

- Deletions of VLAN IP did not previously remove the associated IP records in the HW. This could lead to initial connectivity failure if the deleted VLAN IP address were re-used on a different device until an ARP request was issued by the new device. This issue is resolved (wi00564233).
- Routed ipv6 packets were previously dropped when ECMP was configured for IPv4. This issue is resolved (wi00564341).
- When changing the IP address of a VLAN, the previous VLAN IP address would still be advertised to adjacent devices in the topology tables. This resulted in the "show sys topology" command from the adjacent boxes still displaying the previous IP address of the neighbour until a cold boot of the neighbour was performed. This issue is resolved (wi00507312)

DHCP

- DHCP Replies are flooded by Relay Agent instead of sending unicast DHCP reply packet. This issue is resolved (wi00564847).

IP Multicast

- A PIM message received on a specific port was discarded if the egress towards the source is on the same port but different VLAN. This issue is resolved. (wi00564777)
- In multicast over SMLT environments, aggregation switches could never age out certain S,G entries even after the source stopped sending multicast traffic. These stale entries would eventually prevent new entries from being created and could also result in high cpu utilization. This issue is resolved (wi00697430)

MLT / SMLT

- An infrequent scenario has been identified which could result in a CPU reset in the SMLT module during upgrades to 5.1.2.0. This issue is now resolved. (wi00564781).
- The VLAN MAC addresses of an IST peer switch should be learned only via IST links in a normal operating environment. It has been noted in some transient network scenarios that flooded control plane packets may be reflected and result in these addresses being learnt temporarily on SMLT links. SMLT

behaviour has now been enhanced to ensure that the IST peer switch VLAN MAC addresses are never learnt on SMLT links. (wi00733200).

- In full mesh SMLT with both LACP and VLACP configured, a higher than expected number of VLACP packets may be forwarded to the CPU. This issue is resolved (wi00564193).
- In ERS 8600 with static-mcast mac configured in a SMLT setup, when the links are disabled/enabled, the error message "rarIsPortInMgid(4095): Invalid Mgid" was seen. This issue is resolved (wi00564321).
- When both aggregation switches in an IST pair are rebooted with SMLT ports disabled in the saved configuration file, the SMLT state will remain "normal" on one of the switches when the SMLT ports are subsequently restored. This issue is resolved (wi00564214).
- In an LACP SMLT, the SMLT status would incorrectly remain in SMLT state if the established LACP aggregation times out on the link without a link failure. This issue is resolved (wi00828218)
- In scenarios where packets generated by an IST peer are reflected back on an SMLT link, the "self" mac address associated with IST peer will no longer be learnt on an SMLT link and will remain properly programmed on the IST MLT. This improves resiliency during transient conditions, such as network convergence, that may result in packet(s) being reflected back to the IST peer on an SMLT link. (wi00733200)

BGP/VRF

- Disable/enable of route-policy configured within a VRF results in a crash because of NULL pointer dereference. This issue is resolved (wi00564289).
- On disabling/re-enabling BGP or doing BGP restart on the VRF, routes updates were not being processed properly in configurations where VRF 1 was not configured. This issue is resolved (wi00564348).



Software Release 5.1.3.0

Release Date: August 06th, 2010

Purpose: Software maintenance release to address externally found customer issues.

Changes in This Release

New Features in This Release

None.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- The command “show ports stats show-all” was previously not displaying all of the associated port statistics information. The ERS 8600 now will display the proper information. (Q02012952-01)
- The egress queue service rates configured in NNCLI mode were previously getting lost after a reboot. The egress queue service rates configured are now retained after a reboot. (Q02116358)
- In NNCLI mode “show qos egress-queue-set <queue-set-id>” command was showing no output. ERS 8600 now displays the proper output on executing this command. (Q02116618)
- Previously initiating a traceroute via JDM to an unknown destination while other ping or traceroute activities were being performed simultaneously from the same switch could potentially lead to system instability. This situation has now been addressed. (Q02137566)

Platform

- In specific scenarios with filters and port mirroring being used simultaneously, the ERS8600 will no longer see an increase in CPU utilization due to packets being sent improperly to CPU. (Q02141867)
- In some specific ACL configurations with default-action specified, the ERS8600 previously blocked certain traffic patterns improperly. The situation has now been addressed. (Q02068243)
- Previously links on an R-module could stay up for 60 seconds when the primary SF/CPU (of a dual non-HA SF/CPU configuration only) was not properly removed (module pulled out without prior master reset or switchover). The ERS 8600 now ensures that the links are brought down immediately when the primary SF/CPU is removed, for this configuration. (Q02102654)
- The radius secret key will now be stored as encrypted in the shadv.txt file. Before upgrading to 5.1.3.0, it is now required of the user to delete the radius secret key and then re-add it after the up-grade is complete to avoid accessibility problem with RADIUS. This situation has already been addressed in v7.0 code release. (Q01881817-03)
- When multiple failovers are performed on an ERS 8600 in HA environments, some R module I/O modules could come back up off-line. This situation is now addressed in the 5.1.3.0 code release but also requires the use of the new updated DPC FPGA image (see File Names section). Associated with this 5.1.3.0 Release, the

DPC firmware image must be upgraded for all R-module I/O modules to the dpc194.xsvf firmware image. If the firmware image is not upgraded, the user will receive a warning log message upon any re-boot of 5.1.3.0 code release, warning them that DPC FPGA firmware image is out of revision. (Q02053766-01)

- On E/M module cards, receiving 802.3x pause frames on gigabit Ethernet ports ,could previously result in port level resets. This situation is now addressed. (Q02101087)
- For an ERS 8600 running with 8692 SF/CPU with SuperMezz, the SuperMezz physical LED will now display properly . (Q02102364)
- During boot-up, the potential for 8612XLRS 10Gig port flapping will no longer occur. (Q02138423)
- An error in loading configuration file previously resulted in all of the line cards being disabled. This behavior was added in 5.1 code. This behavior is now changed such that the rest of the correct config will be loaded and all the remaining line cards will not be disabled, but only the line card associated with the improper configuration will be disabled. This situation is different than an “invalid config file, with verify-config flag enabled”, in which case the system will not load the config and will bring up the system with all I/O modules disabled (versus loading the default config). This situation has already been addressed in v7.0 code release. (Q02056382-01)
- A power usage calculation error has been corrected for 8648TXE and 8691SF cards which previously lead to improper warnings being generated in some configurations indicating that the chassis was running on low power. (Q02072016)
- The value of the ingress records was showing improper values via the command show ip mroute-hw resource usage. This has been now been addressed and the proper values for these records are now displayed. (Q02120442)
- The Software power tables embedded in the ERS 8600 were out of sync with the Power Supply Calculator posted via the web. This inconsistency has now been addressed. This situation has already been addressed in v7.0 code release. (Q02020261-01)

RSTP/MSTP

- It had been seen that there was an outage of 30sec in some RSTP setups when one of the root ports was disabled. This situation has now been addressed and the ERS 8600 re-converges within the proper time interval. (Q01984762-02)
- IN RSTP mode, the log file transfer feature was not working properly. This situation has been addressed and the log file transfer feature is now working as expected when RSTP mode is enabled. (Q02101565-01)

IP Unicast

BGP

- In a specific aggregation scenario, when an ERS 8600 forms a neighbour relationship with Cisco, the BGP session will no longer go down due to a malformed AS path. (Q02085844-01)
- During some specific BGP transition scenarios, system instability was previously seen for the ERS 8600. This situation is now addressed. (Q02134841)
- BGP instabilities are no longer observed associated with a HA failover. (Q02135118)
- After an HA failover with a BGP route policy enabled, some BGP routes could be rejected and not re-advertised by the ERS 8600. This situation has been addressed. (Q02136224)

IP Multicast

- ERS 8600 now ensures that when an IP Multicast sender and the receiver are connected to the same ERS 8600, that SPM (Source Path Messages – specific message type within PGM) packets are no longer dropped. This is independent of PGM being enabled or not, as PGM packet flows can function with PIM-SM enabled. (Q02119454)

MLT / SMLT

- ERS8600 will now deterministically (and properly) hash traffic flows for IPVPN traffic on MLT links. (Q02142913)
- MLTs will now come up properly when the LACP Min Link feature is enabled. (Q02034692-02)
- In some rare network events, it is possible to learn the fdb-entry for an IP interface of the peer aggregation in a Switch Cluster on the SMLT associated port instead of the IST_MLT. In this scenario, the fdb-entry will now be properly updated afterward to correctly point to the IST_MLT. (Q02142730)
- CPU high buffer utilization and associated IST instabilities will no longer be observed in specific situations for the ERS 8600. (Q02109963)
- The 8632TXE module will now properly do offline in the scenario where both master and slave CPU's are removed in systems running with smlt-on-single-cp enabled. (Q02123052)

OSPF

- After an upgrade it is now ensured that the OSPF Md5 keys are no longer lost. (Q02127996)

VLACP

- If a mismatching VLACP configuration exists on two ends of a connection associated with a SLT, the SLT will no longer show as SMLT up improperly after a reboot. (Q02119095)

VRRP

- Enabling and then disabling the IST protocol will no longer lead to improper values displayed in the VRRP records. (Q02106080-01)

BFD

- On a ERS8600 running both BGP and BFD, when one tries to enable BGP first and then BFD on the BGP neighbour, the BFD related config can now be set properly. (Q02141063)



Software Release 5.1.2.0

Release Date: April 16, 2010

Purpose: Software maintenance release to address externally found customer issues.

Changes in This Release

New Features in This Release

None.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- The SNMP trap for rclpBgpPeerLastError will now be sent with a proper byte string length such that the last byte will no longer be lost. This could previously cause operational issues with some SNMP management stations. (Q02092718)
- ERS 8600 will no longer observe system instability associated with configuration changes to switch parameters involving SNMP settings. (Q02094258)
- Previously the ERS 8600 was applying a local Access Policy to IPv6 routed SSH packets. Now the system will route these packets and apply Access Policies to only local destination policy type (SSH, Telnet, HTTP) IPv6 packets. This will no longer cause inappropriate connection issues to remote hosts. (Q02070640-01)
- *SNMP GET/GET NEXT under certain conditions associated with the MAC (FDB) table was previously not working properly. These conditions lead to an issue that NetIQ PVQM would not function properly with an ERS 8600. This is now resolved. (Q02113802)*
- ERS8600 has been modified to now allow proper communication with NetQOS Management Device (generally used for IPFix data collection) via SNMPv3. (Q02049612-01)

Platform

- With both filtering and ingress mirroring enabled on the ERS8600, system instability could be seen under certain traffic conditions. This is now resolved. (Q02078239-01)
- IP fix traffic from the switch to an external collector will no longer be sent with an improper QoS marking of QOS=7, but instead sent with QOS=0, now placing these packets into the proper default egress queue. Previously this traffic could potentially interfere with other system management traffic leading to the potential for system instability when IPFix was enabled. (Q02044640-01)
- High CPU utilization on an I/O module co-processor (therefore R/RS only) will no longer result in a loss of messaging synchronization with the 8692 SF/CPU, which previously could have led to system instability. (Q02085085)



- ERS 8600 will no longer show system instability in while writing to the PCMCIA card with CLI Logging enabled. (Q02006689-01)
- ERS 8600 R and RS module card ports will now initialize multicast and broadcast bandwidth limiting values properly when these features are enabled. (Q02074960)
- *Previously when an 8630GBR experienced SW messaging instability to the CP (COP to CP messaging) that would required a module reset to resolve, the far end port could stay up for 20 seconds, which could lead to a 20 second SMLT blackhole. This was associated with the MAC chip not turning off its laser when the module was reset. This situation is now resolved, and the 20 second SMLT blackhole will no longer be seen.* (Q02112285)
- ERS 8600 will now properly handle IPX packets with a broadcast destination MAC of type RIP or SAP. Previously this could create a potential issue for routing IPX for E/M modules (R/RS modules do not support IPX Routing). (Q01997486-04)
- Packet throughput performance for jumbo frames at line rate has been improved for the 8612XLRS modules. (Q02075673)
- Filter pattern definitions for HTTP packet streams will no longer impact other protocol traffic. (Q02089688)
- Users will now be able to connect to an ERS 8600 using Secure Copy (SCP) with access-level rwa when access-strict true is also configured. Previously SSH worked, but SCP did not. (Q01767930-01)
- ERS 8600 will no longer encounter link flapping upon reboot of an OM1400 edge device running SFFD when connected to 8630GBR ports. Avaya continues to recommend the use of VLACP over SFFD, except in cases where the Avaya (or ex-Nortel) product does not support VLACP, such as the OM1400. (02014236-01)
- ERS8600 will now properly forward DHCP packets with the DHCP-relay agent configured as the VRRP virtual IP when the DHCP request has the broadcast flag set. Avaya best practice recommendation continues to be to configure the DHCP-relay agent IP address as the VLAN physical address and not use the VRRP IP address. (Q02059607-01)
- Reliability of R and RS series line card recovery after CPU resets (normally seen during switch software upgrades) has been improved due to enhancements in SF/CPU to I/O module co-processor message communication and synchronization. (Q02091485/ Q01997485)
- ERS 8600 will no longer silently drop packets when the number of ACEs with debug count enabled is such that system resources are at their maximum, but instead the filters will now all function properly. (Q02045086)
- *A memory corruption scenario has been identified which previously led to intermittent system instability and log messages relating to data manipulations of the rarHashBin data structure in some environments (log messages would contain RAR wording). The underlying cause of the memory corruption leading to the instabilities has now been resolved.* (Q02093533, Q02106560)

RSTP/MSTP

- Enhanced MSTP/RSTP logging information which was previously added in release 4.1.3.0 was not present in any 5.x code. This functionality has now been properly added. Q02053232)
- The VLAN interface on an ERS8600 in RSTP/MSTP mode will no longer be brought up unless a port first becomes active in the VLAN. This matches the existing VLAN interface behaviour in STP mode. (Q02083039)
- Packet loss on an MLT with RSTP enabled will no longer been seen after a CPU reset/switchover with HA mode enabled or after a complete switch re-boot. (Q02003158-01)
- ERS 8600 will properly retain the MLT path-cost configuration over reboots when configured for RSTP/MSTP mode. (Q02048253)

- ERS8600 will now properly show the MSTP CIST port path cost info when "show port info mstp" is executed. (Q02048252)

IP Unicast

UDP

- The configured filter action is now properly observed for ACL's configured to match UDP source and destination port ranges between 32752 and 32767. (Q02076252-01)

Static Routes

- ERS 8600 will no longer encounter system (DRAM) memory exhaustion with DHCP-relay configured on a Layer 2 VLAN or at the port level for a non-brouter port. (Q02076879)

BGP

- ERS 8600 will now properly learn the default routes from eBGP peers even after the failover or toggling of the physical port connection. (Q02094999)

IP Multicast

- ERS 8600 will no longer observe periods of sustained high CPU utilization associated with the forwarding of multicast traffic. (Q02067852)
- ERS 8600 will now properly recover its DVMRP status for an ATM interface when a Port/Fiber Fault occurs, and is then restored. (Q02041428)

MLT / SMLT

- Connectivity to NLB servers single homed to one ERS8600 in an IST pair will now function properly for SMLT connected devices when using an nlb-mode of unicast or with arp multicast-mac-flooding enabled. Configurations using nlb-mode of multicast were not affected. (Q02037778-01)
- *When the lowest member port of an SMLT is in an operationally down state during SMLT recovery, ARP records would previously be programmed incorrectly during the database synchronization and cause connectivity failures until the SMLT port was recovered. ARP records are now properly programmed for this scenario. (Q02124545)*
- SLPP will now disable the correct SMLT port when a loop is detected on an SMLT link where the smlt-id configured is not the same as the mlt-id value configured. (Q02089994)
- On ERS8600, FDB and ARP entries will point correctly to SMLT after IST peer reboots. Previously entries learnt on SMLT ports could very occasionally point incorrectly to the IST. (Q02091486)

RSMLT

- With ICMP redirect enabled on RSMLT peer switches, packets destined to the RSMLT-peer's MAC address will now be forwarded correctly and not dropped as ICMP-redirect packets. (Q02091034)
- In RSMLT environments, ERS8600 will no longer add the RSMLT-peer's MAC address to its Router MAC table. This will result in packets destined to the IP interface of RSMLT-peer to forward properly. (Q02091350)

SLPP

- *Previously SLPP would fail to detect a loop if the port uptime was between 24.8 and 49.7 days. After 49.7 days, and then again at 24.8 days, the same situation would repeat - the faulty state would cycle between working to non-working, back to working, etc. based upon the days time period. This is now resolved. A work-around to this situation previously is either a switch reset/re-boot, or port disable/enable prior to 24.8 days of operation. (Q02113609)*
- For non-routed VLANs, SLPP will now use a source MAC address equal to the Base Mac Address of the ERS8600 plus the ID of the VLAN. This will ensure that received SLPP packets are processed against the correct non-routed VLAN when a loop is present in the network and avoid erroneous warning messages. (Q02081719)

VLACP

- ERS 8600 will now always bring down a port via VLACP within the configured timeout value when its VLACP peer goes down. Previously one end of the link would take an extra timeout cycle before downing the port in some scenarios. (Q02088710)
- In scenarios where a port was taken down by VLACP and then the far end switch is rebooted or VLACP recovered to recover the port, Persistent VLACP port flapping will no longer occur. (Q02088709)
- On E-mode enabled switches in full mesh SMLT topologies, protocol traffic will now flow properly on the second MLT link when the first MLT link is disabled. (Q02089615)

VRRP

- Disabling and re-enabling the IST session on an IST switch pair with VRRP configured between them will no longer result in both switches reporting VRRP master ownership. (Q02104773)



Software Release 5.1.1.1

Summary

Release Date: October 2009

Purpose: Software maintenance release to address software issues found both in the field and internally.

Changes in This Release

New Features in This Release

None.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- Network reachability testing via ICMP/ping will no longer show different results when used via either Out of Band connectivity (console or OOB Ethernet port) or when used via an Inband (telnet/SSH, etc.) connection, as was previously seen. (Q02057984-01)

Platform

- Previously certain IST message handling could be delayed by other system functions, thereby potentially causing IST instability (up/down/up). As well, system instability associated with SMLT (IST Peers) in association with high CPU utilization and potentially SLPP operations have both now been resolved. For those who disabled SLPP on their systems/network, SLPP can now be re-enabled with the 5.1.1.1 release. (Q02055292-02/Q02053200/Q02055101/Q02066500)

MSTP

- For an MSTP enabled system, port disable and enable scenarios, where the cistforceport state is disabled on the port, will now be handled properly and in turn OSPF will behave normally. (Q02064812)

IP Unicast

Static Routes

- Static routes usage in a VRF configured system for non-default VRFs (non-VRF 0 usage) will no longer cause a spike in CPU utilization. Now even after reboot all the static routes will remain active and CPU utilization will remain normal. This situation was introduced in 5.1.1.0 code, so only applies to that specific release. (Q02060978-03)

MLT / SMLT

- Previously unicast traffic could be flooded with the VLAN when some SMLT/RSMLT associated link failed; this is now resolved. (Q02037171)



Software Release 5.1.1.0

Release Date: 7 August 2009

Purpose: Software maintenance release to address software issues found both in the field and internally.

Changes in This Release

New Features in This Release

With this release, ERS 8600 introduces new commands to better handle receiving bad OSPF LSAs. The Switch will have an option to configure the way the router behaves on receiving a bad LSA. There are now different options on how to handle a received BAD LSA (with hole in mask). This can affect how adjacency is form to other routers in the network. (Q01997413)

The following commands have been implemented for this new functionality:

```
config ip ospf bad-lsa-ignore <enable|disable>
```

To enable the Switch to keep accepting the bad LSAs (with hole in mask) use the following CLI command (default behavior is disable):

```
config ip ospf bad-lsa-ignore enable
```

Alternatively use the following NNCLI command:

```
bad-lsa-ignore enable
```

Other associated NNCLI commands would be:

```
no bad-lsa-ignore [enable]  
default bad-lsa-ignore [enable]
```

Setting the ospf bad-lsa-ignore parameter to enabled maybe required to maintain adjacency with other non-Nortel switch/routers, especially Cisco models.

The same commands under VRF configuration mode are for CLI:

```
config ip vrf <vrf-id> ospf bad-lsa-ignore <enable|disable>
```


and for NNCLI:

```
ip ospf bad-lsa-ignore enable
```

as well as:

```
no ip ospf bad-lsa-ignore enable  
default ip ospf bad-lsa-ignore [enable]
```

To execute these commands OSPF needs to be disabled globally first.
There is no JDM support for these commands at this time.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- ERS 8600 will no longer experience VRRP transitions, ping failures or potential OSPF slowdown or failure when there is binary transfer of a file via FTP or TFTP with a file size greater than the free memory available on the flash. ERS 8600 could previously experience these issues while its CPU utilization was high. (Q01978884-02)
- Switch will no longer show system instability on quitting from a SSH session, even if a SSH File Transfer Window is opened from the existing SSH session more than once. However the ERS 8600 still does not support any File Transfers from a SSH session. (Q01856195-03)
- ERS 8600 no longer allows adding a route in net mgmt table, if the same route already exists in the normal routing table (due to some routable VLAN or static configuration). (Q01987429-02)

Platform

- ERS 8600 no longer allows (a guard rail has been added) port mirroring from Legacy Port to GTR/GTRS port in “Rx mode” and “both mode”. An invalid port number or failed

message will be returned to the user. This operation is allowed for Tx mode only. (Q01790729-02)

- ERS 8600 no longer experiences unexpected Mezz CPU failover with the Mezz card enabled and then saving the configuration via JDM. (Q01981161-01)
- The potential for traffic interruption associated with the Gig ports on the 8634XGRS module has now been resolved. (Q02010160-03)
- ERS 8600 console will now no longer spool repeated messages of AA1419049-E6 (LX) SFP insertion after a switch reboot. (Q01940440-02)
- ERS 8600 will now be able to properly detect all versions of the AA1419049-E6 SFP (1000Base-LX) even after any switch reboot. (Q01980528-02)
- Link flap detect feature is now supported for R-modules cards. (Q01783494)

IP Unicast

RIP

- For ERS 8600, the set metric parameter in a route-policy will now take effect for RIP. (Q01959361-02)

BFD

- On an ERS 8600 running a BFD session over a static route, when a BFD failure occurs the static route will no longer get learned in the routing table, even when the ARP for next-hop is present until the BFD session gets re-established again. (Q02010174)

BGP

- Operational problems with BGP software that could have led to system instability issues have been resolved. (Q02026274 and Q01972590)

OSPF

- OSPF routes will no longer get improperly deleted even while routes are getting added with ECMP enabled. (Q02021239)

MLT / SMLT

- In RSMLT edge support enabled mode, the creation of a new RSMLT enabled VLAN interface **only on one aggregation box** will no longer cause the static default route to get deleted from the hardware, and thereby affect RSMLT forwarding and re-convergence time. (Q02005454-02)
- In a dual SF/CPU configuration, when the last CP card is pulled out or fails, the RS I/O modules will now have their entire ports drop link automatically. This will help in any SMLT designed network to provide better and faster recovery. (Q01991517-02)
- The ERS 8600 will now check for the SMLT status of an MLT only if it is configured as an SMLT, while sending a MAC-address-learn message for any MAC learnt on the MLT. If the MLT is not configured for SMLT, then SMLT status will not be checked and an MAC-address-learn message will be sent to the IST peer. (Q02036964)
- 8600 will now update the ARP when a MAC learn message is received from IST peer, irrespective of whether the MAC is already existing as local or not. This reduces the chances of improper forwarding in SMLT/RSMLT designed networks. (Q02044582)

Multicast Routing Protocol

PIM

- ERS 8600 will now properly forward packets to the DR when the egress port to the DR is the same as the incoming port and the port to the DR changes for some reason. (Q01907611-04)

Platform

- The change for CR 1767930 which is related to proper operation of SCP (Secure Copy) while using Access Policies, which was fix in the 4.1.x stream back in 4.1.6.3, is missing from all 5.x streams. This will be resolved in all future 5.x code streams, but is still missing in 5.1.1.0.

Configuration

- While configuring ds-field under “config ip traffic-filter filter <filter-id> match”, if we give the six dscp bits, it is taking the command improperly. (Q02056382)



MLT

- An MLT with LACP enabled along with min-link configured, may not have ports added to the MLT properly. (Q02034692)

STATIC ROUTE

- When a static route is created within the non-default VRF, it may not become active and high number of such static routes may lead to increased CPU utilization. For system running with only the default VRF (VRF 0) this is of no concern. For those running with multiple VRFs, use of static routes in the non-default VRF (outside of VRF 0) should be limited or not used at all; instead use some routing protocol, such as OSPF or RIP. (Q02060978)

Copyright © 2011 Avaya Inc - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark are trademarks of Nortel. The Ethernet Routing Switch 8100/8300/8600 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 Technical Support on the web hosted by Nortel: <http://www.nortel.com/support>