




Ethernet Routing Switch 8600
Readme
Software Release 4.1.8.5



Table of Content

Ethernet Routing Switch 8600.....	1
Readme	1
Software Release 4.1.8.5	1
Table of Content	2
Software Release 4.1.8.5	4
Important Notices	4
Platforms Supported	5
Notes for Upgrade.....	5
File Names for This Release.....	5
Version of Previous Release	7
Compatibility	7
Changes in This Release.....	8
New Features in This Release.....	8
Old Features Removed From This Release	8
Problems Resolved in This Release.....	8
Outstanding Issues	8
Known Limitations	8
Documentation Corrections	8



Release Date: September 30, 2010

Purpose: Software maintenance release to address a PCMCIA driver issue.

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)

In prior releases, the SNMP timer task could potentially crash if the SNMP retry counter is set to a value greater than zero and multiple concurrent SNMP inform events are generated which do not receive an acknowledgement. While multiple factors and conditions need to align in order to encounter this SNMP task crash, it was recommended to set the SNMP retry count within the snmp-v3 target-address to zero in order to avoid the issue altogether. An example of such a configuration is (see the bolded entry):

```
Config snmp-v3 target-addr create "NNM" 10.10.10.1:162 "TparamV2" tdomain ipv4_tdomain
timeout 1500 retry 0 taglist informTag mms 484
```

This is now resolved in this release, and the retry value can now be set to values greater than zero. (Q02052753)

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 dependant 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 modules

NOTE: R-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 4.1.8.3 (present in this document)

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.nortel.com/support.

NOTE: If upgrading to 4.1.8.5 code within any SMLT designed network from a release prior to 4.1.8.2, 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.

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_4185.tar.gz	53026571
Ethernet Routing Switch images			
Boot monitor image	CPU and switch fabric firmware	p80b4185.img	1091005
Run-time image	Run-time image	p80a4185.img	9148708
Run-time image for R modules	Image for R modules	p80j4185.dld	1272740
Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz)	Image for the SuperMezz card	p80m4185.img	9260261
3DES	Encryption module for privacy protocol with Secure Shell (SSH)	p80c4185.img	55928
AES	Encryption module for privacy protocol for SNMPv3. Includes AES and DES	p80c4185.aes (this image includes the DES image)	26112
MIB	MIB files	p80a4185.mib	3330132
MIB (zip file)	Zip file containing MIBs	p80a4185.mib.zip	533453
MD5 checksum file	md5 checksums of all Release 5.1 software files	p80a4185.md5	1442
Runtime image for ATM	Runtime image for the ATM module	p80t4185.dld	906024
Runtime image for POS	Runtime image for the POS module	p80p4185.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	dpc184.xsvf	2583454
PIM8630GBR	Programmable I/O module FPGA firmware; for the 8630GBR only	PI_769.xsvf	2284578
SSL images			
SSL cluster upgrade	Ethernet Routing Switch 8600 clustered SSL modules self-installing runtime image/upgrade	p80s4185.pkg	5988896
SSL boot monitor	Ethernet Routing Switch 8600 SSL module boot monitor	p80s4185.img	7528448
SSL upgrade instructions	Ethernet Routing Switch 8600 SSL upgrade instructions	p80s4185.upgrade	1481
SSL installation instructions	Ethernet Routing Switch 8600 SSL installation instructions	p80s4185.install	2895
SSL diagnostics	Ethernet Routing Switch 8600 SSL diagnostics	p80s4185.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.i mg	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_6090_solaris_sparc. sh	181208765
Microsoft Windows image	Device Manager software image	jdm_6090.exe	170755822
Linux image	Device Manager software image	jdm_6090_linux.sh	176359101
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.lo0	variable

Version of Previous Release

Software Version **4.1.8.3**

Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.

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

- 8692 CP boards with the new PCMCIA controller could experience high cpu utilization and other systemic effects while accessing certain types of PCMCIA cards. The problem has now been addressed. (wi00564790)

Outstanding Issues

Please refer to the Outstanding Issues Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.8.3. No other new issues have been found.

Known Limitations

Please refer to the Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.8.3. No other new known limitations have been found.

Documentation Corrections

None.

Copyright © 2010 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>



Ethernet Routing Switch 8600

Software Release 4.1.8.3

Software Release 4.1.8.3.....	9
Software Release 4.1.8.3.....	10
Software Release 4.1.8.2.....	14
Software Release 4.1.7.2.....	27
Software Release 4.1.7.1.....	31
Software Release 4.1.6.0.....	41
Software Release 4.1.5.4.....	52
Software Release 4.1.4.0.....	66
Software Release 4.1.3.0.....	76
Software Release 4.1.2.0.....	92
Software Release 4.1.1.1.....	113



Ethernet Routing Switch 8600 Software Release 4.1.8.3

1. Release Summary

Release Date: 22 May, 2009

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

2. Important Notes before Upgrading to This Release

Important Note: Users should read and reference CSB 2008008618, Rev 3, Software Life-Cycle Management for the ERS 8600 product, before deciding to move to any code release. An alternative code consideration to Release 4.1.8.3 may be Release 5.0.1.0. Release 5.0.1.0 is available as a GA web posted release. All CRs addressed in release 4.1.8.3 are also addressed in release 5.0.1.0, and release 5.1.0 also includes the SMLT re-architecture changes that were incorporated into release 4.1.8.3.

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.8.2 for information on upgrades with SMLT. Customer upgrading from 4.1.8.2 to 4.1.8.3 can use normal upgrade procedures for switch cluster upgrades. The "Important Notes" section in the 4.1.8.2 portion of this document applies only to customers upgrading their switch clusters to 4.1.8.3 or higher from a software release prior to 4.1.8.2 (4.1.7.2 and below).

3. 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 dependant on the presence of the High Performance Backplane. Reference 4.1.1.0 RN document below for further details.

The following modules are not supported in the 8003 chassis:

- 8692SF/CPU
- 8630GBR
- 8648GTR
- 8683XLR
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)



- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Name	Module or File Type	File Size (bytes)
p80b4183.img	Boot monitor image	1090914
p80a4183.img	Runtime image	9148468
p80j4183.dld	Run-time image for R modules	1272740
p80c4183.img	3DES	55928
p80c4183.aes	AES (this image includes the DES image)	26112
p80a4183.mib	MIB	3330132
p80a4183.mib.zip	MIB (zip file)	533453
p80a4183.md5	md5 checksum file	1442
p80p4183.dld	8600 POS module image	701771
p80t4183.dld	8600 ATM module image	906024
p80m4183.img	8600 image for the SuperMezz card	9258801
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4183.pkg	SSL cluster upgrade	5988896
p80s4183.img	SSL boot monitor	7528448



p80s4183.upgrade	SSL upgrade instructions	1481
p80s4183.install	SSL installation instructions	2895
p80s4183.diag	SSL diagnostics	19460381

WSM Images

wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004

SDM Firewall Images

NSF5100_2.3.7.0_SDM_R65.img	NSF Boot image
NSF5100_2.3.7.0_SDM_R65.iso	NSF Boot ISO
NSF5100_2.3.7.0_SDM_R65.pkg	NSF upgrade package

SDM TPS Images

Nortel_3D_Sensor_SDM-4[1].7.0.2-Restore.iso	Upgrade script (patch) to upgrade TPS from 4.5.x to 4.7.0.2
---	---

5. Version of Previous Release

Software Version **4.1.8.2**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.

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.

7. Changes in This Release

Old Features Removed From This Release

None.

AVAYA

Problems Resolved in This Release

Switch management

- Repeated SSH sessions, with CLI command logging enabled, could potentially lock up and no longer allow any other connections to the node. JDM or console access could be still available initially and further lost as well once attempting to run any output command. This problem could lead to a complete Switch Fabric card software lockup. (Q02001015)

Platform

- The 10GE ports on an 8683XLR card may flap for traffic rates at or very close to line rate. (Q01860167)

SMLT/Layer 2 switching

- There can be a discrepancy in the ARP tables between IST peers where one peer has invalid or duplicate ARP entries, ARP entries with a negative TTL value or ARP entries which do not appear to be aging. (Q01999818)
- After an IST MLT is flapped or the IST Peer is rebooted, invalid ARP entries can get created on the IST Peer for a MAC that is present in multiple VLANs. (Q01992234).
- ARP entries are found to be persistently learned on the IST-MLT link while the FDB entries for the associated MAC addresses were properly learned on the SMLT links when multiple MLT/SMLTs returned to service at about the same time causing an overlap in IST synchronization for those MLTs/SMLTs. (Q01996928)

8. Outstanding Issues

The issues listed in the 4.1.8.2 "Outstanding Issues" section of this document all apply. In addition, the following issue has been observed and has not been corrected in this release.

- Q02028752: SSH related Error messages observed

Error messages "Unable to fetch the source address" and "SW ERROR slot2src: Can't get ssh ip address" may appear on the console while logging in continuously via SSH for a few hours. A data access exception message may also appear on the console or in the log with CLI logging enabled. It has been determined that these error messages and the data exception are not service impacting to the operation of the ERS 8600, nor do they prevent SSH sessions from being established.

9. Known Limitations

Same as Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.8.2.

10. Documentation Corrections

Same as Documentation Corrections Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.8.2.



Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600 Software Release 4.1.8.2

1. Release Summary

Release Date: 30 January, 2009

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

NOTE: For those customers who were part of the 4.1.8.0 Controlled Release program, an upgrade to 4.1.8.2 is not required. The CRs listed in *italic* below are the external CRs changed between 4.1.8.0 and 4.1.8.2.

2. Important Notes before Upgrading to This Release

Important Note: Users should read and reference CSB 2008008618, Rev 3, Software Life-Cycle Management for the ERS 8600 product, before deciding to move to any code release. An alternative code consideration to Release 4.1.8.2 maybe Release 5.0.1.0. Release 5.0.1.0 should be available as a GA web posted release by the end of February, 2009.

When upgrading a pair of IST Core peer switches, special care must be taken due to SMLT architectural change (see Section 7 - New Features). Unlike previous upgrades, when moving from pre- 4.1.8.x code to any 4.1.8.x or higher code, the IST peer switch operation will be different during the upgrade process. When the first peer switch is upgraded (secondary SSF/CPU first, and then primary SSF/CPU), this switch will go into an ALL port lock state, with the IST_MLT being the only links/ports up and active. The SMLT connected portion of the network will continue to operate on the second IST Core switch, still running the pre-4.1.8.x image. The user should now wait approximately 5 minutes until the following message is seen on the upgrade 4.1.8.x switch:

CPU5 [05/15/08 05:33:55] MLT ERROR SMLT initial table sync is delayed or failed. Please check the peer switch for any partial config errors. All the ports in the switch except IST will remain locked

At this time, the second IST Core peer switch can be upgraded. At that time, and when the 1st now running 4.1.8.x switch sees its IST_MLT connection go down, it will unlock all its ports, and start learning and then forwarding traffic. This upgrade process will cause some user traffic impact, even for the SMLT connected



portion of the network. The 2nd IST Core switch once finished upgrading should show the following messages in its log – these should be checked for to confirm proper upgrade and operation:

```
CPU5 [06/05/08 05:05:45] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take ANY action with the peer at this time
CPU5 [06/05/08 05:05:46] MLT INFO SMLT MAC /ARP Sync is Complete : Peer can now be used normally
```

After both IST Core peer switches are upgraded and network is running 'normally', it is now recommended to power-cycle (re-boot only for single SSF/CPU systems) the one peer switch which was upgraded first so that a complete "sync" (see Section 7) will now occur.

After upgrading to 4.1.8.x or higher code, the future upgrade process will return to the behavior previously documented. When one IST Core peer switch is upgraded, the SMLT portion of the network will continue operation without user impact on the other IST Core peer. The second IST Core peer switch can then be upgraded in the same manner, resulting in little to no user traffic impact for the SMLT connected portion of the network.

This new upgrade operation affects SMLT based networks moving from pre-4.1.8.x code to 4.1.8.x or higher code. If moving from 4.1.8.x to 5.0 based code, it is recommended to ONLY move to 5.0.1.0 or higher code, as this code has similar SMLT enhancements and therefore a 'normal' IST Cluster upgrade can be done. Upgrades from 4.1.8.x to pre-5.0.1.0 code will NOT be supported.

NOTE: After upgrading there is a very small potential for invalid (but not operationally affecting) ARP entries to show up in the ARP table. These entries start with an IP address with the number 0 (0.x.x.x). Running the command '*show ip arp info 0.*' will show if such entries exist. If no values are returned, then you are all set. If entries are returned, please contact Nortel Support and request Level 2 GNPS assistance.

3. 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 dependant on the presence of the High Performance Backplane. Reference 4.1.1.0 RN document below for further details.

The following modules are not supported in the 8003 chassis:

- 8692SF/CPU
- 8630GBR
- 8648GTR
- 8683XLR
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade



Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4182.img	Boot monitor image	1087545
p80a4182.img	Runtime image	8798967
p80j4182.dld	Run-time image for R modules	1272524
p80c4182.img	3DES	55928
p80c4182.aes	AES (this image includes the DES image)	26112
p80a4182.mib	MIB	3329553
p80a4182.mib.zip	MIB (zip file)	533374
p80a4182.md5	md5 checksum file	1768
p80p4182.dld	8600 POS module image	701771
p80t4182.dld	8600 ATM module image	906024
p80m4182.img	8600 image for the SuperMezz card	8909615
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4182.pkg	SSL cluster upgrade	5988896
p80s4182.img	SSL boot monitor	7528448
p80s4182.upgrade	SSL upgrade instructions	1481
p80s4182.install	SSL installation instructions	2895
p80s4182.diag	SSL diagnostics	19460381
<i>WSM Images</i>		



wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004

SDM Firewall Images

NSF5100_2.3.7.0_SDM_R65.img	NSF Boot image
NSF5100_2.3.7.0_SDM_R65.iso	NSF Boot ISO
NSF5100_2.3.7.0_SDM_R65.pkg	NSF upgrade package

SDM TPS Images

Nortel_3D_Sensor_SDM-4[1].7.0.2-Restore.iso	Upgrade script (patch) to upgrade TPS from 4.5.x to 4.7.0.2
---	---

5. Version of Previous Release

Software Version **4.1.7.2**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.

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.

7. Changes in This Release

New Features in This Release

SMLT/RSMLT Operational Improvements (CR Q01764193/Q01769324/Q01776485)

For previous SMLT operation, bringing the SMLTs Up/Down triggered the flushing of the entire MAC/FDB belonging to the SMLTs in both the IST Core Peer Switches. Flushing of the MAC addresses then causes the dependent ARP (for IP stations) to be re-resolved. For ARP resolution, ERS 8600 re-ARPs for all the SMLT learned ARPs. This created a major MAC/ARP re-learning effort. As the records were flushed, during the re-learning period the exception (learning) packets will also be continuously forwarded to the CPU, thereby increasing the CPU load. This would further slow-down the SMLT re-convergence as well as the h/w record re-programming. Since proper traffic flow with an ERS 8600 is completely dependent on the h/w records, this prior behavior could adversely affect convergence times, especially in very large networks (8000+ MACs/ARPs), and those networks also running with many multicast streams, as multicast streams often need to be forwarded to the



CPU for learning, thereby also increasing CPU load.

The SMLT changes in this release improve this operation significantly, and continue to allow all previous SMLT/RSMLT topologies to be supported. SMLT Operational Improvements will affect SMLT/RSMLT behavior in that the actual SMLT/RSMLT connection links on a powered-up IST Core Switch, will take longer to become active (link status up and forwarding) than with previous versions of code. During this time period the other Peer IST Core Switch will always be continuing to forward, therefore avoiding any loss of traffic in the network for all SMLT/RSMLT based connections. The SMLT/RSMLT associated links will not become active upon a boot-up until the IST is completely up, and a 'new MAC/ARP/IP sync' has occurred between the two Core IST Peer Switches in a Cluster.

Users may see occasional instances where the Remote SMLT Flag is False on both Peer Switches. This is normal, if the flag clears and is then set properly (False on one side, True on the other), once the FDB age-out for that associated VLAN has occurred. This behavior has no affect on user traffic operation – no user traffic loss or disruption will be seen under this condition.

For proper network behavior Nortel recommends to operate both IST switches with either the "new" or "old" SMLT architecture. Therefore SMLT operation between IST Peer Core switches with one switching operating with pre-4.1.8.x code, and the other operating with 4.1.8.x or later code is NOT supported.

Additionally users will see some new informational log messages generated around this behavior. The new messages formats are listed below, along with the various situations they will be seen with.

Case 1: Switch running SMLT is reset. Upon switch coming up the below messages are displayed irrespective of the number of SMLTs:

```
CPU5 [06/05/08 05:05:45] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take ANY action with the peer at this time
CPU5 [06/05/08 05:05:46] MLT INFO SMLT MAC /ARP Sync is Complete : Peer can now be used normally
```

Case 2: System is up and running but SMLT UP event (from down) has happened. One sync message is displayed for every SMLT that went down and has come up. In the following example, 2 x SMLTs went down and came up:

```
CPU5 [06/05/08 05:05:45] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take ANY action with the peer at this time
CPU5 [06/05/08 05:05:45] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take ANY action with the peer at this time
CPU5 [06/05/08 05:05:46] MLT INFO SMLT MAC /ARP Sync is Complete : Peer can now be used normally
CPU5 [06/05/08 05:05:46] MLT INFO SMLT MAC /ARP Sync is Complete : Peer can now be used normally
```

NOTE: To determine which specific SMLT IDs are affect, look for the SMLT ID down/up log messages.

Case 3: When sync fails due to difference in IST Peer software version (pre-4.1.8.x and 4.1.8.x) where one peer supports MAC/ARP sync but the other does not. Or some other potential issue, such as a mis-configuration or IST Session not coming up. The system that is reset and is requesting sync, it will keep all the ports locked down (except IST_MLT) until the IST comes up properly and sync has occurred. After 5 minutes the below Log/Error messages will be displayed:

```
CPU5 [05/15/08 05:28:51] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take ANY action with the peer at this time
< After 5 min>
CPU5 [05/15/08 05:33:55] MLT ERROR SMLT initial table sync is delayed or failed. Please check the peer switch for any partial config errors. All the ports in the switch except IST will remain locked.
```

NOTE: All known failover times for SMLT/RSMLT operation are now, and always have been sub-second. With this release all known fail-back or recovery times have been improved, especially for very large scaled environments to be within 3 seconds, in order to provided required redundancy for converged networks. These



values are for unicast traffic only. Not all IP Multicast failover or fail-back/recovery situations can provide such times today, as many situations depend on the IPMC protocol recovery. For best IPMC recovery in SMLT/RSMLT designs, the use of static RPs for PIM-SM is recommended, with the same CLIP IP address assigned to both Core IST Peers within the Cluster, and to all switches with a full-mesh or square configuration. Failover or fail-back/recovery times for any situations that involve high-layer protocols can not always be guaranteed. Reference the Network Design Guide for your specific code release for recommendations on best practices to achieve best results. In many situations, it is abnormal corner case events for which times are extended. As well for all best results, VLACP MUST also be used. The SMLT/RSMLT improvements noted here have been optimized to function always with VLACP. Therefore for best results a pure Nortel SMLT/RSMLT design is best. We still support SMLT designs with any non-Nortel devices that support some level of link aggregation, but fail-back/recovery times can not be guaranteed.

NOTE: VLACP configuration should now use values of 500 msec short timer (or higher) and a minimum timeout-scale of 5. Lower values can be used, but should any VLACP 'flapping' occur, the user will need to increased one or more of the values. These timers have been proven to work for any large scaled environments (12,000 MACs), and also provide the 3 second recovery time required for converged networks (5 x 500 = 2.5 seconds). Using these values may not increase re-convergence or fail-back/recovery times, but instead guarantee these times under all extreme conditions. (CR Q01925738-01 and Q01928607) As well, users should note that if VLACP is admin disabled on one side of the link/connection, this will cause VLACP to bring remote connection down, but since remote connection will keep link up, the side with VLACP admin disabled, will now have a black-hole connection to the remote switch, which will drop all packets being sent to it. If VLACP is disabled on one side of a connection, it MUST also be disabled on remote side or else traffic loss will likely occur.

NOTE: If using VRRP with SMLT, users are now HIGHLY (MUST) recommended to use unique VRIDs, especially when scaling VRRP (more than 40 instances). Use of a single VRID for all instances is supported within the standard, but when such a configuration is used in scaled SMLT designs, instability could be seen. A [better] alternative method, which allows scaling to maximum number of IP VLANs, is to use RSMLT designs instead. See Section 10 in this Readme (page 10) for additional information on how to easily move from VRRP design to RSMLT design.

NOTE: For any SMLT design, for L2 SMLT VLANs, it is now HIGHLY recommended to change the default VLAN FDB aging timer from its default value of 300 seconds, to now be 1 second higher than the system setting for the ARP aging timer. FDB timers are set on a per VLAN basis. If using the default system ARP aging time, *config ip arp aging <minutes>*, of 360 (minutes) than the proper value for the FDB aging timer, *config vlan x fdb-entry aging-time <seconds>*, should be 21601 seconds, which is 360 minutes (6 hours) plus 1 second. This will have the system only use the ARP aging timer for aging, versus the FDB aging timer. This value has been shown to work very well to assure no improper SMLT learning. The use of this timer has one potential side-affect. For legacy module, this limits the system to around a maximum of 12,000 concurrent MACs; for R-mode system, the limit remains at 64K, even with timer setting. With this timer, should an edge device move, the system will still immediately re-learn and re-populate the FDB table properly, and not have to wait for the 6 hour (plus 1 second) timer to expire. No negative operational affects are known when using this timer value. For non-SMLT based VLANs the default FDB aging timer of 300 maybe used or can be changed or even also set to 21601. For this reason the default value of the FDB aging timer will remain at 300 (seconds), within all code releases.

DDMI SFP Support (Q01785224)

The following DDMI SFPs are now supported:

Part Number	Part Description
AA1419048-E6	1-port 1000Base-SX Small Form Factor Pluggable (SFP) Gigabit Ethernet Transceiver, connector type: LC. Digital Diagnostic Monitoring Interface
AA1419049-E6	1-port 1000Base-LX Small Form Factor Pluggable (SFP) Gigabit Ethernet Transceiver, connector type: LC. Digital Diagnostic Monitoring Interface
AA1419050-E6	1-port 1000BaseXD Small Form-factor Pluggable (SFP) Gigabit Ethernet Transceiver - 1310nm. Diagnostic Monitoring Interface
AA1419051-E6	1-port 1000BaseXD Small Form-Factor Pluggable (SFP) Gigabit Ethernet Transceiver - 1550nm. Digital Diagnostic Monitoring Interface



AA1419052-E6	1-port 1000BaseZX Small Form-Factor Pluggable (SFP) Gigabit Ethernet Transceiver 1550nm. Digital Diagnostic Monitoring Interface
AA1419053-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1470nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419054-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1490nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419055-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1510nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419056-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1530nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419057-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1550nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419058-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1570nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419059-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1590nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419060-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1610nm Wavelength, 40km. Diagnostic Monitoring Interface
AA1419061-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1470nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419062-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1490nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419063-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1510nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419064-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1530nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419065-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1550nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419066-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1570nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419067-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1590nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419068-E6	1-port 1000BaseCWDM Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1610nm Wavelength, 70km. Diagnostic Monitoring Interface
AA1419069-E5	1-port 1000Base-BX Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1310nm Wavelength. Must be paired with AA1419070
AA1419069-E6	1-port 1000Base-BX Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1310nm Wavelength. Must be paired with AA1419070
AA1419070-E5	1-port 1000Base-BX Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1490nm Wavelength Must be paired with AA1419069
AA1419070-E6	1-port 1000Base-BX Small Form Factor Pluggable GBIC (mini-GBIC, connector type: LC) - 1490nm Wavelength. Must be paired with AA1419069

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- Users will now be allowed to enable ssh-pass-auth when the switch is operating in Hsecure mode. (Q01818955)
- Users will now be allowed to access a switch operating in Hsecure mode with JDM after setting the block-snmp bootconfig flag to false. This value is not currently preserved over a reboot – see the Outstanding Issues section below for more information. (Q01811342)



- Removal of non-operational I/O cards will now be detected by ERS 8600 and both the CLI and JDM will properly display the card as not present. (Q01862981/ Q01855789)
- When hard or soft resets are now initiated from JDM or other SNMP devices, a log message indicating the reset type will now be generated. These log messages will have format:

CPU5 [01/30/09 16:53:19] SNMP INFO SWITCH EXTERNALLY REBOOTED

(Q01983973)

- Disconnecting an SSH session that had initiated a secure file transfer (an unsupported protocol) can no longer lead to system instability. (Q01856195)
- Modification of the switch login password through JDM will now be communicated to the WSM card properly and users will be able to access the WSM module after the password change. (Q01867943)
- If the verify-config flag is set to true, ERS 8600 no longer defaults the configuration on boot if an R module fails to initialize or is not present, even though there exists configuration information for that slot. (Q01865292)

Platform

- Port statistics for 8683XLR modules no longer reflect incorrect high values after a port is disabled and re-enabled. (Q01890883)
- Some versions of a new chassis, any chassis with HW Rev field of 50 or above, where coming out of manufacturing with the HW Config field set with invalid value. This type of chassis would the show invalid log message errors regarding Slots 1 and 10 (for 8010 or 8010CO chassis only) not being High Performance slots, when indeed the slots were HP slots. Loading 4.1.8.2 or higher code into such a chassis will correctly reprogram the affect field and the chassis will no longer report the invalid error log messages. (Q01950243-01)
- ERS8600 no longer mirrors traffic to the mirroring port when the configuration related to port mirroring is disabled or deleted. (Q01849168)
- ACE filters that match on an IP protocol and have subsequent ACE filters that match on source or destination IP will now function properly after the switch is rebooted. (Q01814530)
- ERS8600 operating in super-mezz mode will now recognize the PCMCIA card properly when the PCMCIA card is removed and re-inserted without issuing a PCMCIA-stop command before removal, as well as after initial insertion of the PCMCIA card when the switch was booted without a PCMCIA card. (Q01907447)
- CPU modules are no longer polled for parity errors, as the poll result inaccurately reflected an error condition. This operation is associated with the detection of FAD / SWIP errors (ER Q01505098 – see 4.1.6.0 Readme). (Q01910296)
- A request for statistics on a port which is part of a VLAN on which an ACL is applied and when that port or a port in the same slot is not specified in the ACE attributes will no longer cause switch connectivity issues. (Q01865206)
- Ping snoop filters will now be properly excluded from the configuration file when a save config is performed. (Q01895651)
- A rare lockup of a data path "lane" on a 8630GBR line card that can occur during a card or switch reboot has been resolved. This lockup would only occur occasionally when custom queues were configured. (Q01777369)



- The ambient temperature parameter for new CMHS fans shall now be displayed properly for a system running on release 4.1.x code stream. Release 4.1.8.2 or higher fully supports the new CMHS fans. These fans now ship with any new chassis, not just for chassis where RS modules are to be deployed. (Q01938086)

Layer 2 switching

- FDB entries will no longer incorrectly display the SMLT Remote flag as False/False or True/True on both the IST pair switches for the same MAC entry. (Q01838871)
- Ports RSTP/MSTP related configuration will now be retained after a removal and re-insertion of an I/O module. (Q01783263/Q01978551)
- MLT ports running RSTP no longer experience traffic forwarding issues when individual MLT ports go up or down. (Q01944595-03, Q01940344)
- Legacy module SMLT ports containing 2 or more links configured with RSTP will now properly come up following a reboot. (Q01943168)
- A CPU failover event will no longer incorrectly cause VLACP to bring SMLT ports down. (Q01894048)
- ERS8600 no longer displays an error message when a link is brought down which has a static-mcastmac configured. (Q01881405)
- In SMLT environments, rebooting one of the IST peers or toggling the status of SMLT no longer can cause system instability; this change is associated with SMLT re-architecture. (Q01908107)
- VLACP will now converge properly on SMLT ports after HA failover as port locking and unlocking is now synced properly to the slave CPU. (Q01913866)
- MSTP will now properly send BPDUs on HA failover and will no longer cause instability in the network. (Q01887792)
- MSTP will now converge properly when it contains ports that are not part of the CIST instance. (Q01921984)

IP Unicast

- ERS8600 now forwards traffic ingressing on an MLT port when Reverse-Path-Check is enabled in strict mode. (Q01859755)
- The R module error message “COP-SW ERROR Slot #: ercdSetDefaultRoute: rcdRspMalloc failed” that was being generated during default route updates has been resolved and no longer causes system instability. (Q01908628-02)
- ERS8600 RSMLT edge support now recovers properly when last active port in RSMLT VLAN is bounced. (Q01953415-02)

BGP



- When a route-policy is applied to match a particular network and a BGP neighbor is restarted, the "show ip bgp neighbor advertise-route x.x.x.x" command will now display only the routes that are advertised to the BGP Peer. (Q01879357)
- After deleting a BGP aggregate address, BGP will no longer incorrectly advertise routes back to the originating BGP peer. (Q01889887)
- When an AS-List filter is disabled and an IBGP peer is restarted, IBGP will now advertise BGP routes properly. (Q01889914-03)

Multicast

- An ERS8600 with an R-module will no longer become unstable when excessive multicast control packets are received on that module. (Q01872811)
- On an ERS8600 using PIM-SM, the multicast route entries will no longer point to the wrong next hop when there is a change in the network topology. (Q01907611)
- A method to solve the issue of connection Id leakage for IP Multicast has been provided, This resolves the issue of failed attempts in creating a connection to the MGID when creating S,G records, and thereby potentially affecting IP Multicast operation. (Q01906890)

8. Outstanding Issues

- Enabling of the multicast flooding feature (config ip arp multicast-mac-flooding enable) does not sync to the slave when the ERS8600 is operating in HA mode. (Q01942511)
- SNMP/JDM connectivity is lost after ERS8600 is reset with hsecure flag is set to true (enabled) and block-snmpp flag set to false. To correct, go back and set the block-snmpp flag to false. In a future release, hsecure operation will be modified. (Q01811343/ Q01901801)
- The QoS value of unicast packets is retained when forwarded to the CP as exception packets. If enough packets with high QoS setting are received, this could negatively affect CP handling of other packets. In general, unicast packets being sent to CP is abnormal, and the root cause of that situation should be investigated and resolved as a first step. (Q01845219)
- On an ERS running BGP and OSPF, when BGP routes are redistributed into the OSPF domain and a route-policy is used to match and permit a prefix, more specific prefixes do not get redistributed into the OSPF domain. Care must be taken when using such a configuration. (Q01922909)
- VRRP global configuration parameters are set to disable (default is enabled) on deleting a configured prefix list. This will not affect VRRP operation, but will affect ability to ping the VRRP IP address, and for VRRP trap operation. This situation only happens upon a prefix list delete, and requires a reconfiguration of VRRP global parameters to currently resolve. (Q01924298)
- Configuring Distributed MLT on ERS8600 using 8608GBE cards could lead to a ports being put in different STP states in the MLT. (Q01877552)
- The "show sys info card" CLI command executed for an 8692 CPU card with Mezz Hardware version E2 does not display the proper Mezz Serial number. (Q01921632)
- When an 8648GTR module is connected to classic ethernet modules through an MLT, only the first MLT port added auto negotiates correctly to 100Mbps Full Duplex whereas the subsequent ports added comes up as 10Mbps Half duplex. (Q01886904)



- On an ERS8600 running OSPF, the command “show vlan info ports [vlan-id]” display inadequate space between the column head VLAN ID and PORT MEMBER when the ports to be displayed are OSPF passive port members. (Q01927867)
- If UDP forwarding is configured on IST switches, the destination subnet will receive two copies of the original packet, as both IST switches will receive and forward the packet. If the forwarding address is a subnet broadcast and that VLAN is configured on both of the IST switches, instability may be observed on the IST link. (Q01936551)
- A PIM BSR may experience instability when a large number of RP sets are configured. For large PIM implementations, it is recommended to use summarization for RP groups and not use /32 masks. A warning message will now be generated under these circumstances, stating: "BSR message size is high at XXXX bytes. Summarize your RP Sets to reduce their size. Failure to do this could lead to system instability.". This message will be displayed when the BSR message size is 1719 bytes or greater. This situation will not occur in networks with static RP configurations where no BSR is present. (Q01938405)
- When MSTP or RSTP is configured, an SNMP get-next on any MIBs in the tree: 1.3.6.1.2.1.17.1.4.1 (iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1). dot1dBridge(17). dot1dBase(1). dot1dBasePortTable(4).dot1dBasePortEntry(1)) will return the same OID as the request, potentially resulting in high CPU utilization. (Q01942480).
- The PCMCIA card may become inaccessible and the following error may be shown:
0x53e7ab0 (tLoggerTask): disk cache error: device fe21488 block 57 errnoc0003, disk write failed, data loss may have occurred.
This error is most likely caused by the switch not running the latest DOSFS code. Workaround: 1) backup the files from the DOSFS device 2) reformat the DOSFS device with Release 5.0 or later (use the dos-format command) 3) reinstall the files on the DOSFS device. (Q01986665)

9. Known Limitations

- Broadcast and multicast rate limiting may occasionally cause packets to be dropped when traffic levels are below the configured rate limit when configured on legacy modules. (Q01871916).
- The event of discontinuing logging to PCMCIA by the executing the CLI command pcmcia-stop is not logged into the log file. (Q01774929).
- When VLACP is configured, the log message “LACP WARNING Received MAC-mismatched PDUs on port <port number>, MAC <mac address>, Please check your VLACP configuration.” is sometimes observed. In the event of a VLACP misconfiguration, this message will be logged continuously in the log file, and the configuration should be corrected. If VLACP is not misconfigured and the message is observed sporadically, it should be disregarded as there is no impact to traffic. (Q01938504)
- ERS8600 could take a minimum of about 10 seconds to populate an OSPF learned route into the Route Table Manager as it does not rebuild its router LSAs immediately when the neighbor router changes to FULL state. (Q01848350)
- A “save config” executed from the standby CPU will result in the pcap.cfg data replacing the configuration data stored in config.cfg on the standby CPU. Changes to the standby CPU configuration should ONLY be executed from the master CPU, via either “save config” (savetostandby flag set true or “save config standby <config file name>”. (Q01937023)
- SLT configurations where the SLT port and the IST port are configured on the same physical module (non-best practice design) can experience FDB entry learning issues for MAC addresses learned across



the IST, which can cause connectivity loss. This condition can occur if the module is physically swapped out during switch operation, or if the module is disabled/enabled via CLI or JDM. All FDB entries will be relearned after the fdb-ageout time, or can be manually relearned by performing a fdb-entry flush on the relevant VLANs. (Q01926665)

- On an ERS8600 running in SMLT environment using E modules, if the number of ARP and FDB table entries exceeds 7500, the ARP table entries may experience intermittent corruption when the corresponding FDB table entries age out. (Q01904966/Q01869054)
- On an ERS8600 running VLACP, if the value of VLACP timer is not set to a multiple of 10, VLACP may not function properly because the VLACP port may not be able to send out any VLACP PDUs. VLACP timers must always be set to some multiple of 10. (Q01932414)
- When multiple ATM PVCs are configured on a port and a static route is configured where the ATM interface is the next hop, the static route will not be removed from the routing table when a PVC goes down as detected by F5-OAM. The static route will only be removed when the ARP entry ages out. For a single PVC configuration, the static route will be removed as the ATM port will go down. (Q01952948)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

The CLI command "config sys set reset-passwd snmp-community-strings" is now obsolete and no longer valid. The documentation in Configuration and Managing Security for ERS8600 regarding this command is no longer valid. (Q01923596)

Please note the following in regards to mixed module (R/RS and classic E/M) MLT support:

Mixed module MLTs consisting of R and RS module ports and classic E and M module ports will work together and is fully supported only for common features that function for both types of modules. Many functions and parameters must be set at the individual port level within the MLT. VLAN assignment, tagging (802.1d) status, and Spanning Tree Protocol (STP) are examples of functions that are configured at the MLT level and not at the port level. If a parameter is only available on a subset of the MLT ports, than the use of this parameter is not supported in a mixed module MLT, and with such use inconsistent behavior can occur. Speed and duplex setting are examples of port level parameters that are supported on all types of modules, and therefore can always be used. Note that common port level parameters or functions must always be configured to be the same for all ports within the MLT. Multicast and broadcast rate limiting is a feature available on both port types, but it operates differently depending on the module type. Therefore, Nortel recommends not to use this function, as it could lead to inconsistent or unknown behavior. Egress queue size is a parameter that is only available with R and RS modules and should not to be used in a mixed module MLT, as the parameter cannot be applied to the E and M module ports. You need to use caution when using filters in a mixed module MLT. Filter operation can differ between the module types. If you need to use a filter, you need to ensure the different configurations create filters that apply the same rules and actions for both port types.

There is little to no software control or checking for mixed module MLT configuration. Before configuring a mixed module MLT, you should address any questions or concerns to Nortel prior to implementation.

Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.



To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600

Software Release 4.1.7.2

1. Release Summary

Release Date: 19 November 2008

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

2. Important Notes before Upgrading to This Release

None.

3. 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 dependant on the presence of the High Performance Backplane. Reference 4.1.1.0 RN document below for further details.

The following modules are not supported in the 8003 chassis:

- 8692SF/CPU
- 8630GBR
- 8648GTR
- 8683XLR
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 31545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.



File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4172.img	Boot monitor image	1073723
p80a4172.img	Runtime image	8774693
p80j4172.dld	Run-time image for R modules	1272856
p80c4172.img	3DES	55928
p80c4172.aes	AES (this image includes the DES image)	26112
p80a4172.mib	MIB	3329497
p80a4172.mib.zip	MIB (zip file)	533346
p80a4172.md5	md5 checksum file	1442
p80p4172.dld	8600 POS module image	701771
p80t4172.dld	8600 ATM module image	906024
p80m4172.img	8600 image for the SuperMezz card	8873804
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4172.pkg	SSL cluster upgrade	5988896
p80s4172.img	SSL boot monitor	7528448
p80s4172.upgrade	SSL upgrade instructions	1481
p80s4172.install	SSL installation instructions	2895
p80s4172.diag	SSL diagnostics	19460381
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004



NSF5100_2.3.3.0_SDM_R60.img NSF Boot image

NSF5100_2.3.3.0_SDM_R60.iso NSF Boot ISO

NSF5100_2.3.3.0_SDM_R60.pkg NSF upgrade package

SDM TPS Images

Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso TPS Intrusion Sensor Boot image

Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso TPS Defense Center Boot ISO

Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1

Nortel_TPS_IS_Upgrade_XXX_Upgrade-10.md5 IS upgrade download verification file.

Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1

Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5 DC upgrade download verification file

5. Version of Previous Release

Software Version **4.1.7.1**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.

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 2.3.7.0. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

This software release supports SDM TPS Release 4.7. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

To download these code releases may require valid Nortel support web access.

7. Changes in This Release

New Features in This Release

None.



Old Features Removed From This Release

None.

Problems Resolved in This Release

IP Unicast

- An ARP processing situation, introduced in the 4.1.6.0 release, that could lead to system instability has been resolved in this release. Customers experiencing system instability with code releases starting with 4.1.6.0 and up to and including 4.1.7.1 should move to this code release. Customers may want to consider a pro-active approach of moving to 4.1.7.2 depending upon their current network operation and code release. This change is the only change between 4.1.7.2 and 4.1.7.1. (Q01923361, Q01925711)

8. Outstanding Issues

Refer to the 4.1.7.1 Release Notes below.

9. Known Limitations

Refer to the 4.1.7.1 Release Notes below.

10. Documentation Corrections

- The CLI command "config sys set reset-passwd snmp-community-strings" is now obsolete and no longer valid. The documentation in Configuration and Managing Security for ERS8600 regarding this command is no longer valid. (Q01923596)

Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600

Software Release 4.1.7.1

1. Release Summary

Release Date: 02 Sep 2008

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

2. Important Notes before Upgrading to This Release

Release 4.1.7.0 was never made into a GA release due to CRs Q01913866/Q01894048-03. These CRs dealt with VLACP operation not functioning 100% in HA-CPU enabled configurations. 4.1.7.1 solves both of the above CRs (same underlying situation for both), and therefore 4.1.7.1 will replace 4.1.7.0 as the GA release. Reference the Issues Fixed section for additional details.

3. 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 dependant on the presence of the High Performance Backplane. Reference 4.1.1.0 RN document below for further details.

The following modules are not supported in the 8003 chassis:

- 8692SF/CPU
- 8630GBR
- 8648GTR
- 8683XLR
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.



File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4171.img	Boot monitor image	1073755
p80a4171.img	Runtime image	8773803
p80j4171.dld	Run-time image for R modules	1272856
p80c4171.img	3DES	55928
p80c4171.aes	AES (this image includes the DES image)	26112
p80a4171.mib	MIB	3329497
p80a4171.mib.zip	MIB (zip file)	533346
p80a4171.md5	md5 checksum file	1442
p80p4171.dld	8600 POS module image	701771
p80t4171.dld	8600 ATM module image	906024
p80m4171.img	8600 image for the SuperMezz card	8873615
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4171.pkg	SSL cluster upgrade	5988896
p80s4171.img	SSL boot monitor	7528448
p80s4171.upgrade	SSL upgrade instructions	1481
p80s4171.install	SSL installation instructions	2895
p80s4171.diag	SSL diagnostics	19460381
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256



wsm1003400_boot.img

WebOS boot image

43004

SDM Firewall Images

NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package

SDM TPS Images

Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1
Nortel_TPS_IS_Upgrade_XXX_Upgrade-10.md5	IS upgrade download verification file.
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh	Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5	DC upgrade download verification file

5. Version of Previous Release

Software Version **4.1.6.0/4.1.6.3**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.

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 2.3.7.0. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

This software release supports SDM TPS Release 4.7. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

To download these code releases may require valid Nortel support web access.

7. Changes in This Release

New Features in This Release

None.



Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- If a DNS server is configured on the switch and a user does a cut/paste of commands into a telnet session, all the commands within the cut/paste will now be executed properly. (Q01784202)

Platform

- Garbage characters ingressing the console port at a high rate of speed will no longer result in system instability. (Q01825851)
- When a SuperMezz card is enabled and excessive traffic destined to an unknown host is sent to an in-band VLAN, and there is also a condition of continuous SNMP requests to the switch, the switch will no longer become unstable. (Q01802463)
- A manufacturing issue was identified on some 8005AC power supplies in which the power supply was incorrectly programmed. This led to all prior versions of software to not properly recognize these power supplies. This recognition issue would cause the power supply output power to be defaulted to 690W in system power calculations, as the power supply will show as 'not recognized'. Other than this the power supply will function 100% properly. This lower output power calculation could potentially lead to modules not coming online in a highly populated chassis. This release of software corrects this issue and properly recognizes the power supply and applies the correct output power to the system power calculations. (Q01920749)
- For Classic modules, if both enable-diffserv and access-diffserv are enabled on a mirrored port, users will no longer be able to create a port mirror with mode Tx or Both, just Rx. Previously this mirror configuration was allowed but resulted in potentially corrupted mirrored packets. (Q01744991)
- SMLT links configured with VLACP will no longer be intermittently disabled by VLACP after an HA failover. This situation was seen only after a chassis power cycle, or reset of both SSF/CPU's at the same time, AND then a follow-on HA failover. If your system gets into this situation, with an older release of code, a global disable/enable of vlacp (config vlacp dis, then config vlacp en) will rectify the situation. In order to prevent future failures, after any system/chassis power cycle a reset of the stand-by SSF/CPU, would be recommended, and will then alleviate the issue from returning. At the same time, because this only occurs with SMLT ports, even if this situation did occur the SMLT peer will continue to stay up and run the network without interruption. (Q01913866/Q01894048-03)
- On ERS 8600 when port mirroring is configured on R modules, the Tx frames will now be mirrored properly for mirroring mode Tx or Both. (Q01842897)
- When a faulty secondary SSF/CPU is present in a chassis, the Master SSF/CPU will no longer continuously attempt to bring the faulty SSF/CPU online, potentially causing high CPU utilization. (Q01869443-03)

Layer 2 switching

- For Classic fiber modules, the SMLT status will no longer be incorrectly displayed as "smlt" for MLTs which are down. (Q01867402)



- Performing a CPU failover on one peer switch within and SMLT pair no longer causes system instability on the other IST Peer switch due to improper handling within the ARP delete functionality. (Q01907052)

IP Unicast

- When viewing the IP -> Policy -> ApplyPolicy tab for the ERS 8600 via Device Manager, the fields RoutePolicyApply, RedistributeApply, and OSPFInFilterApply are now unchecked by default. This now matches pre-4.1.x behavior. There is no reason to have every Policy type checked by default, as the user needs to select (check) the type of policy they wish to apply by using the apply button associated with this screen. Users should be aware that 4.1.7.1 and above will not function differently than previous 4.1.x releases in how this DM screen functions. Please also note the 4.1 based IP Routing manual will not be updated and therefore page 639 shows the pre-4.1.7.1 behavior, not the new 4.1.7.1 behavior. This change has no other functional or operational change. (Q01834790)
- When a DHCP/BootP server is connected to a VLAN that is in VRRP backup state and a client is connected to a VLAN on a different 8600 in VRRP master state, the DHCP/BootP replies will now be forwarded. Previously in such a configuration the packets were only forwarded by the switch in VRRP backup state when VRRP backup master was also enabled. (Q01848171)
- ERS8600 in an SMLT environment will no longer become unstable with certain traffic patterns also involving a large number of ARP records. This situation is detailed in CSB 2008008980, which was also fixed in 4.1.6.3 code. The 4.1.7.1 release has this 4.1.6.3 change included. (Q01862845)

OSPF

- When ERS 8600 receives an invalid LSA, it will now send out a BadLSReq which will bring down the neighbor relationship with the router sending the invalid LSA. (Q01837492)
- OSPF routes with a next hop corresponding to the IP address of the RSMLT Peer will no longer be removed from the route table if the IST_MLT is toggled. (Q01894652)

BGP

- In a BGP configuration with a hold-timer of less than 5 seconds and a keepalive-timer of less than 2 seconds, insertion of an R module will no longer cause a loss of BGP peers. (Q01825956)
- BGP input policies are now properly applied when a soft-restart is performed. (Q01939163)

Multicast

- If an IP prefix list is created via JDM or CLI with a name that contains spaces, and it is used for configuring an IGMP access control entry, a reboot will no longer cause the configuration to be lost. (Q01829843)
- When multiple IGMP access control lists are configured on a brouter port, the configuration is no longer lost upon switch reset. (Q01829919)
- For PIM over SLT (single link SMLT) configuration, multicast traffic is no longer lost when one of the SLT links is disconnected. (Q01836125)



- ERS 8600 is no longer occasionally unstable when multicast mode NLB is configured. This situation is explained in CSB 2008008980. (Q01885431)
- For SMLT triangle configurations in which the edge box is running IPMC (PIM-SM), the IST peer switch with the lower IP address will no longer stop forwarding traffic downstream upon receiving traffic from the upstream edge router. (Q01890745)

IPX

- ERS 8600 will no longer route IPX packets with a broadcast destination MAC unless the packet type is RIP/SAP. (Q01749446)

8. Outstanding Issues

- The CLI command "show vlan info fdb-entry" can display MAC entries where the SMLT remote parameter indicates either both true or both false for the same MAC entry on each peer IST switch. This condition has not been seen to create any traffic passing issues. (Q01838871)
- When port mirroring is configured on a classic port with RxFilter mode and an ingress filter with an action of mirror is also applied to that same port, packets are still mirrored when port mirroring is disabled. The workaround is to disable both port mirroring and the ingress filter action of mirror. (Q01849168)
- When attempting to perform a secure file transfer (an unsupported protocol) from within an SSH session the 8600 will be unstable after leaving the SSH session. Users should not perform unsupported secure file transfers. (Q01856195)
- Module information is incorrectly displayed in CLI and Device Manager after a faulty R module has been removed from the chassis. (Q01855789, Q01862981)
- If Reverse Path Check is enabled in strict mode for a VLAN or Brouter port, traffic ingressing on an MLT will not properly be forwarded. For now RPC should not be used if the VLAN or Brouter port is associated with an MLT. This applies only to R-modules as RPC is an R-module only feature. (Q01859755)
- For 8630GBR modules, port mirroring does not mirror Tx traffic if the 3DES file is loaded. (Q01842007)
- If the login password is changed via Device Manager, the new password will not authenticate on the WSM card. The workaround is to change the password via the CLI. (Q01867943)
- Displaying ACL/ACE statistics via CLI or JDM for an advanced filter that is applied on a VLAN with ports on different physical modules can sometimes result in high CPU utilization, and slow network management responsiveness. (Q01865206)
- OSPF may lose adjacencies when a file is transferred to /flash memory via the management port and there is not enough space within the /flash memory. (Q01870315)
- If the verify-config flag is set to true, the ERS 8600 defaults the config (verify fails) on boot if an R module fails to initialize or is pulled out. To not see this behavior, verify-config flag may need to be set to false (disabled). (Q01865292)
- In a RSMLT set-up where the peer core switches both have a Mezz card, there is a loss of connectivity if the core switch which matches the real default gateway IP address is reset. This situation is not seen in configurations with no Mezz card present. (Q01864544)
- Under conditions of continuous ARP broadcast with large amounts (7500+) of unique MAC addresses are sent to ERS 8600 at a high rate of (200+ packets per second), some ARP entries may temporarily point to



the IST versus SMLT when the FDB entry ages out, resulting in traffic loss. A potential solution to this situation is to set the FDB age-out on a per VLAN basis to 21601 for every VLAN. This value is 1 second higher than the default ARP age-out (6 hours or 21600 seconds), and therefore FDB age-outs will not occur, only ARP age-outs. (Q01869054)

- A poorly worded error message is displayed when attempting to configure the port speed to 1000M with auto-neg disabled on an 8648 GTR module. This setting is not allowed per IEEE 802.3ab standard. In a future release the message will be improved. (Q01849911)
- If a save config to standby fails due to a problem accessing its flash, CPU messages, such as OSPF hellos, may not be sent during the save attempt. (Q01827936)

9. Known Limitations

- Enabling trace level 9 with option "screen on" for a console session may cause IST/OSPF/VRRP transitions under heavy CPU traffic conditions. This is not observed when the "screen on" option is not enabled. Tracing with screen option on is never recommended in production networks. (Q01852005)
- When an I/O module is inserted/removed from a chassis, VRRP transitions may be seen, if the VRRP fast advertisement interval has a value of 400ms or below. These transitions can be ignored and acknowledged after the I/O maintenance activity is completed. (Q01838155)
- ERS 8600 may take around 10 seconds to populate an OSPF route into the routing table after it forms full adjacency with its neighbors. This is the behavior for all code streams, including all older streams, such as 3.x. (Q01848350)
- On R modules a mismatch in the duplex settings may cause improper traffic flows. Users are warned against improper auto-negotiation settings, which can lead to this behavior. Reference CSB 2008008723. (Q01804243)
- During HA failover running MSTP with default timers, BPDUs are not sent for a long enough period of time that other MSTP switches can become root bridge, resulting in a loop in the network. This issue is not seen if the hello timer value is set to minimum of 6 sec and MaxAge to a minimum of 20. (Q01875995)
- In an SMLT set-up with classic modules in non-M mode, if more than 8000 MAC and 8000 ARP entries are sent to a core switch, there may be a possibility of the SMLT peer ARP entry getting removed from the ARP table, as the forwarding table could be full. This situation needs many more addresses to be seen in either M-mode or R-mode. (Q01894730)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

- Although the ERS8600 supports 802.1w (RSTP), the product does not completely conform to the final standard, only the original draft standard. (Q01875995)



APPENDIX

MIB changes in 4.1.7.0 release. (Q01842838-01; Clean-up of RapidCity.MIB)

rcPortIndex OBJECT-TYPE
SYNTAX InterfaceIndex
MAX-ACCESS read-only
STATUS current
DESCRIPTION "An index value that uniquely identifies a port.
This value is similar to ifIndex in MIB2."
 ::= { rcPortEntry 1 }

rcPortType OBJECT-TYPE
SYNTAX INTEGER {
 other(0), -- no port installed
 rc100BaseTX(1), -- (cat 5)
 rc100BaseT2(2), -- (cat 3)
 rc100BaseF(3), -- (100 mbps fiber)
 rc1000BaseF(4), -- (1000 mbps fiber)
 rc1000BaseDualF(5), -- (fiber - dual connector)
 rc10BaseF(6), -- (10 mbps fiber)
 rcPosMMF(7), -- (multi mode fiber)
 rcPosSMF(8), -- (single mode fiber)
 rcAtm(9), -- (oc3, oc12, ds3, e3)
 rcGbicLx(10), -- (gigabit long haul)
 rcGbicSx(11), -- (gigabit short haul)
 rcGbicXd(12), -- (gigabit extended distance)
 rcGbicCu(13), -- (gigabit copper)
 rcGbicOther(14), -- (gigabit other)
 rcOc3cSM(15), -- (OC-3c SM)
 rcOc3cMM(16), -- (OC-3c MM)
 rcOc3cCOPPER(17), -- (OC-3c COPPER)
 rcOc12cSM(18), -- (OC-12c SM)
 rcOc12cMM(19), -- (OC-12c MM)
 rcDs3(20), -- (Ds3)
 rcE3(21), -- (E3)
 rcGbicNone(22), -- (Gbic card feature)
 rc1000BaseT(23), -- (1000 base copper)
 rcGbicZx(24), -- (gigabit very extended distance)
 rcOc3cAtmSM(25), -- (Atm OC-3c SM)
 rcOc3cAtmMM(26), -- (Atm OC-3c MM)
 rcOc12cAtmSM(27), -- (Atm OC-12c SM)
 rcOc12cAtmMM(28), -- (Atm OC-12c MM)
 rcOc3cPosSM(29), -- (Pos OC-3c SM)
 rcOc3cPosMM(30), -- (Pos OC-3c MM)
 rcOc12cPosSM(31), -- (Pos OC-12c SM)
 rcOc12cPosMM(32), -- (Pos OC-12c MM)
 rcGbic1470(33), -- (gigabit wavelength 1470)
 rcGbic1490(34), -- (gigabit wavelength 1490)
 rcGbic1510(35), -- (gigabit wavelength 1510)
 rcGbic1530(36), -- (gigabit wavelength 1530)



rcGbic1550(37), -- (gigabit wavelength 1550)
rcGbic1570(38), -- (gigabit wavelength 1570)
rcGbic1590(39), -- (gigabit wavelength 1590)
rcGbic1610(40), -- (gigabit wavelength 1610)
rcRmon(41), -- (1000 base TProbe)
rcGbic1470APD(42), -- (gigabit wavelength 1470-APD)
rcGbic1490APD(43), -- (gigabit wavelength 1490-APD)
rcGbic1510APD(44), -- (gigabit wavelength 1510-APD)
rcGbic1530APD(45), -- (gigabit wavelength 1530-APD)
rcGbic1550APD(46), -- (gigabit wavelength 1550-APD)
rcGbic1570APD(47), -- (gigabit wavelength 1570-APD)
rcGbic1590APD(48), -- (gigabit wavelength 1590-APD)
rcGbic1610APD(49), -- (gigabit wavelength 1610-APD)
rc10GbLW(50), -- (10 Gig Ethernet LW)
rc10GbLR(51), -- (10 Gig Ethernet LR)
-- OM2.0 place holder added following
-- rcLogicalServerPort(52),
-- opm Prism used as a server for MPLS
rc1000BaseTX(53), -- (triple speed)
rcGbicBx(55), -- (SFP - Single Fiber Bi-Directional - 100Base-Bxl)
rc10GbNone(56), -- (10 Gig Feature)
rc10GbSR(58), -- (10 Gig SR Ethernet)
rc10GbSW(59), -- (10 Gib SW wavelength 850nm and sonet)
rc10GbER(60), -- (10 Gig ER wavelen 1550nm)
rc10GbEW(61), -- (10 Gig EW)
rc10GbOther(62), -- (10 Gig Ethernet other)
rc1000BaseTXPOE(63), -- (triplespeed withPOE support)
rc10GbZR(64), -- (10 Gig ZR wavelength 1550nm)
rc10GbZW(65), -- (10 Gig ZW wavelength 1550nm)
rcGbic1310Xd(66), -- (gigabit wavelength 1310-distance 40Kms)
rcGbic1470Xd(67), -- (gigabit wavelength 1470-distance 40Kms)
rcGbic1490Xd(68), -- (gigabit wavelength 1490-distance 40Kms)
rcGbic1510Xd(69), -- (gigabit wavelength 1510-distance 40Kms)
rcGbic1530Xd(70), -- (gigabit wavelength 1530-distance 40Kms)
rcGbic1550Xd(71), -- (gigabit wavelength 1550-distance 40Kms)
rcGbic1570Xd(72), -- (gigabit wavelength 1570-distance 40Kms)
rcGbic1590Xd(73), -- (gigabit wavelength 1590-distance 40Kms)
rcGbic1610Xd(74), -- (gigabit wavelength 1610-distance 40Kms)
rcGbic1470Zx(75), -- (gigabit wavelength 1470-distance 70Kms)
rcGbic1490Zx(76), -- (gigabit wavelength 1490-distance 70Kms)
rcGbic1510Zx(77), -- (gigabit wavelength 1510-distance 70Kms)
rcGbic1530Zx(78), -- (gigabit wavelength 1530-distance 70Kms)
rcGbic1550Zx(79), -- (gigabit wavelength 1550-distance 70Kms)
rcGbic1570Zx(80), -- (gigabit wavelength 1570-distance 70Kms)
rcGbic1590Zx(81), -- (gigabit wavelength 1590-distance 70Kms)
rcGbic1610Zx(82), -- (gigabit wavelength 1610-distance 70Kms)
rcGbic1690Zx(83), -- (gigabit wavelength 1690-distance 70Kms)
rcGbic1310Bx(84), -- (gigabit wavelength 1310-distance 10Kms)
rcGbic1490Bx(85), -- (gigabit wavelength 1490-distance 10Kms)
rcGbicEx(86), -- (gigabit wavelength 1550-distance 120 Kms)
rcGbic850Sx(87), -- (gigabit short haul wavelength 850)
rcGbic1200Lx(88), -- (gigabit long haul wavelength 1200)
rcGbic1300Lx(89), -- (gigabit long haul wavelength 1300)
rcGbic1310Lx(90), -- (gigabit long haul wavelength 1310)
rcGbic1490Lx(91), -- (gigabit long haul wavelength 1490)
rcGbic1550Lx(92), -- (gigabit long haul wavelength 1550)
rcGbic1550Ex(93), -- (gigabit wavelength 1550-distance 120Kms)
rc1GbFD(94), -- (100/1000 Fiber Gbic)
rc10GbLRM(95), -- (10 Gig LRM)



```
rc10GbDWDMR(96), -- (10 Gig DWDM Ethernet)
rc10GbDWDMW(97)  -- (10 Gig DWDM Sonet)
```

```
}
```

```
MAX-ACCESS read-only
STATUS      current
DESCRIPTION "Port type"
::= { rcPortEntry 2 }
```

Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600

Software Release 4.1.6.0

1. Release Summary

Release Date: 26 March, 2008

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

2. Important Notes before Upgrading to This Release

The new VLACP Global parameters have NOT been added into this release. The addition of these new parameters is still scheduled for a future release.

In a chassis with all R-modules (most likely R-mode enabled) the following flags will not have no affect as these parameters are specific to legacy modules and therefore should always be set to false or disabled:

Control-record-optimization (config bootconfig flags control-record-optimization <false|true>)

Enhanced-operational-mode (EOM) (config sys set flags enhanced-operational-mode <false|true>)

It is also suggested that for best operation that these flags are disabled (set to false) in any mixed chassis that has R-modules present.

The smltRemote log message introduced with release 4.1.5.4, has been changed into a trace message and will not be generated with this code release. Therefore the use of msg-control function to suppress these messages is no longer needed. (Q01745278)

3. Platforms Supported

All Ethernet Routing Switch 8600 modules are supported in the 8006, 8010, and 8010co Chassis. Full slot support for all modules maybe dependant on the presence of the High Performance Backplane. Reference 4.1.1.0 RN document below for further details.

The following modules are not supported in the 8003 chassis:

- 8692SSF/CPU (which means the following R-modules are not supported, however v4.1.x software is support within the 8003 Chassis via the 8690 or 8691SSF/CPU models)
- 8630GBR
- 8648GTR
- 8683XLR
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:



- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are now supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4160.img	Boot monitor image	1073588
p80a4160.img	Runtime image	8772884
p80j4160.dld	Run-time image for R modules	1272848
p80c4160.img	3DES	55928
p80c4160.aes	AES (this image includes the DES image)	26112
p80a4160.mib	MIB	3328560
p80a4160.mib.zip	MIB (zip file)	533162
p80a4160.md5	md5 checksum file	1489
p80p4160.dld	8600 POS module image	701771
p80t4160.dld	8600 ATM module image	906024
p80m4160.img <i>Firmware images</i>	8600 image for the SuperMezz card	8872035
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578



p80s4160.pkg	SSL cluster upgrade	5988896
p80s4160.img	SSL boot monitor	7528448
p80s4160.upgrade	SSL upgrade instructions	1481
p80s4160.install	SSL installation instructions	2895
p80s4160.diag	SSL diagnostics	19460381

WSM Images

wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004

SDM Firewall Images

NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package

SDM TPS Images

Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1
Nortel_TPS_IS_Upgrade_xxx_Upgrade-10.md5	IS upgrade download verification file.
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh	Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5	DC upgrade download verification file

5. Version of Previous Release

Software Version **4.1.5.4**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.0 or higher.



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 2.3.7.0. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

This software release supports SDM TPS Release 4.7. This code is found on Nortel web site under Security & VPN -> Service Delivery Module 8660.

To download these code releases may require valid Nortel support web access.

7. Changes in This Release

New Features in This Release

Partial Support for LACP Standby Links (ER Q01731929, Q01759055)

Configuration of greater than 8 ports in a LACP enabled MLT (MLT with aggregate enabled; LAG) is now permitted. This is not permitted for a 'normal' MLT, where 8 static ports are still the maximum allowed. For any LAG configuration, 8 of the ports will always attempt to be active, while any additional ports will work as standby links. The standby links will become active ports in the LAG if any active link goes down, thereby providing bandwidth redundancy for the LAG. There is no configuration support on the ERS8600 for standby links with any configuration of 8 or less ports. If the far-end device does support standby link configuration with less than 8 ports, the ERS8600 will also now recognize remote ports that are in standby state, and place its appropriate ports in standby state as well. (Q01767527).

Support for link priority (config eth slot/port lacp port-priority) is NOT supported. Therefore one can not 'force' a port from active to standby, or vice-a verse, via setting port-priority for the individual ports within the LAG.

New Log Messages for Ext-CP Limit Feature (Q01756748-01)

The log messages associated with Ext-CP Limit have been changed. The new messages are:

```
CPU5 [03/12/08 15:48:57] SW INFO SysOp remain congested for 5024ms. Total rxDrop = 2724216, qDepth = 0  
CPU5 [03/12/08 15:48:57] SW INFO Ports from Soft down List are causing congestion. starting port monitoring
```

The congested time (milliseconds) should approximate to the configured system minimum congestion timer (config sys ext-cp-limit min-congestion-time). The default value for this system timer is 3000 milliseconds. Total rxDrop will indicate how many packets have been dropped due to the congestion. Note that at this time, this parameter display may be inaccurate (CR Q01845752). This situation will be addressed in a future release. Right now the exact and actual number of packet drops is not as significant as the indication that it is packet drops causing the congestion, versus abnormal queuing. A non-zero qDepth value would be an indication that excess queuing was the cause of the congestion, as either situation can trigger the congestion mechanism to be run.

The 2nd log message is stating the monitoring of Soft Down ports will now start. It should also state that all ports configured as Hard Down enabled will now be disabled. This message will be enhanced in a future code release to state:

"SW INFO Ports from Hard Down List will now be disabled due to SysOp congestion. System will also now starting port monitoring of the Soft Down List."

In addition, the following new message has been added:



CPU5 [03/17/08 13:12:21] SW INFO Port 4/19 has average utilization of xx% for past t Sec.

where, xx% is the average port utilization for the slot/port in question, and t = the configured port-congestion-time (config sys ext-cp-limit port-congestion-time) configured in seconds. The default system port congestion timer is 5 seconds.

This message will only be displayed for ports that actually get disabled by Ext-CP Limit, and have been configured for Soft Down operation. Therefore the above message will also be associated with an existing message of the format:

CPU5 [03/12/08 15:49:02] SNMP INFO Link Down (4/19) due to ext-cp limit

For any configured Soft Down port to be disabled by Ext-CP Limit, the port utilization must exceed the configured port level value for threshold utilization rate (config eth slot/port ext-cp-limit SoftDown [threshold-util-rate]). The default threshold utilization rate is 50%, and valid parameter values are any utilization value between 1 and 100.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- The definition for 10Gbps is now added in rcPortAdminSpeed Mib. Therefore a MIB query for port operational speed will now correctly show 10000 (Mbps) for a 10G port. (Q01780945)
- Radius accounting now shows correct value of input/output packets for a telnet session to any in-band interface address. The number of output packets is still shown as zero for a telnet session to the management port; this is a HW limitation. Also now SNMP-v2 accounting will properly display Client IP-address, while SNMP-v3 accounting will now display the proper name of SNMP user. (Q01759830)
- In the ERS 8600 topology table the name for ERS4500 series switches will now be properly displayed. (Q01798277)
- Users will now be able to connect to ERS 8600 using Secure Copy (SCP) when an access policy with access-level rwa and access-strict is true are also configured. Previously SSH worked, but SCP did not. (Q01767930)
- Changing the Management Port bootp status to enable or disable from JDM will no longer cause management connectivity to be lost, via the Management Port. (Q01746388)
- The ERS 8600 Management IP now does reply to ping if the ICMP packet is marked with DSCP value of AF11. Previously this did not occur. (Q01782522)

Platform

- When a CPU card is present only in slot 6, the tapmux parity error counters (feature introduced in 4.1.5.4 code – ER 1505098) are no longer incremented for non-faulty cards. (Q01807020)
- Traffic loss could be seen on the mirroring port when the mirroring combination was a legacy port with any ingress (Rx) mirror function and mirroring port was on 8648GTR module. Therefore users will no longer be allowed to configure ingress mirroring (any mirror function with Rx) if the mirrored port is a legacy port and the mirroring port belongs to an 8648GTR module. The solution for ingress/Rx mirroring of legacy ports in a mixed chassis is to also use a legacy port for the output mirroring port. (Q01790729)



- On ERS 8600 issuing a Ctrl-X sequence while in a telnet prompt will no longer cause instability. (Q01752502)
- Single Fiber Fault Detection (SFFD) now works properly when an 8630GBR module is connected to an OM1200/OM1400. This is the only current configuration where SFFD should still be used. In all other configurations, VLACP should be used. The OM series of switches do not support VLACP. (Q01757462)
- When the CPU management port is connected to an Ethernet I/O port and speed and duplex is set to 100 Full on both sides, the management port no longer will lose the speed and duplex setting upon a switch reboot. (Q01737868)
- The link-flap-detect feature is now supported on R-modules just as it has always functioned for legacy modules. (Q01783494)
- With enhance-operational-mode flag enabled and a MLT is configured using any WSM ports, the switch is no longer unstable when the WSM card is pulled out. (Q01815827)
- When using PCAP to capture packets, and a large file is transferred between two hosts, PCAP will now work properly, and for the next attempt PCAP packets will no longer be dropped. (Q01754856)
- When the count of parity errors exceeds the user set threshold the CPU is now reset. (Q01505098-02, Q01727771)
- When the PCMCIA system log file exceeds the configuration limit, a warning message informing the user about this situation will now be logged into the local log file (show log file), the system log file stored on the PCMCIA and the message will be transferred to any syslog server as well. The message format is:

SW Warning Warning: THE ALLOWED LOG FILE SIZE HAS EXCEEDED CONFIGURATION LIMITS.
THE FILE SIZE IS CURRENTLY xxxxxxx BYTES!!!!

Additionally if space remains on the PCMCIA card, a new log file with an incremented extension, for example .001, will be created.

(Q01766905)

Layer 2 switching

- A MAC address, for which an fdb-filter is configured, is now displayed in the Forwarding table with type as "mgmt". (Q01497977-03)
- When RSTP/MSTP mode is enabled, a VLAN will now become unreachable when there are no active ports present in the VLAN. This is now the same behavior as normal STP. (Q01797033)
- With RSTP enabled, the configured path cost of a port which is a member of either a MLT or MLT with LACP (LAG) will no longer be lost (reset to factory default value) when the port is disabled/enabled or upon a switch reset. (Q01772097)
- All IST packets are now sent with proper QoS for p-bits of 7, and are therefore handled properly even when diffserv is enabled on the IST MLT ports. (Q01772936)
- With RSTP enabled, disabling and re-enabling of STP on an RSTP edge-port will no longer cause any connectivity issues. (Q01753578)
- When classic IP traffic filters are created with redirect next-hop option set to a some external interface, the ERS8600 will no longer forward traffic to the changed next-hop external interface after the IP Address of the next-hop device is changed or replaced. In order to get such forwarding, the next-hop filter IP Address will now need to be re-configured. (Q01774902)



- With the CPU utilization being high and a log message being repeated a large number of times the CLI command "show log file tail" will no longer cause switch instability. (Q01780128)
- In an SMLT configuration also using HA, FDB entries are now properly synced on both of the IST Peers after an HA failover. (Q01790052-01)
- For Chassis shipped from year 2003 onwards, the CLI command "show sys info" will now properly display MacAddrCapacity as 4096k. This additional capacity, also obtained via the MAC Address Upgrade Kit (DS1411015) for older chassis, is required for VLAN scaling. (Q01723144)
- ERS 8600 now no longer sends out SLPP packets tagged with ID of VLANs which have been deleted but were not removed from the SLPP configuration. (Q01782455)
- The "unknown-mac-discard activation enable" command should be executed as the last command of all port security commands. Otherwise the "unknown-mac-discard" function could be activated with default parameters. In this instance the ERS8600 port will discard all unknown mac packets, even if autolearn has been set enabled. This can create communication issues. Normally this is not an issue, except some users like to copy the output of the "show config" command and paste it into another device's console when they setup another device. Previously the "show config" always display "unknown-mac-discard autolearn enable" after "unknown-mac-discard activation enable", and therefore this caused activation enable to occur before auto-learn could occur. Now the order of lines of the "show config" display has the "unknown-mac-discard activation enable" command displayed as the last line of port security commands. (Q01733255)

MLT/SMLT

- In a SMLT set-up where wireless APs are connected off of edge switches, a situation for which the wireless APs could lose synchronization and reset continuously has now been resolved. This situation was created by the network design, and some of the SMLT changes introduced in the 4.1.3.0 release, as this situation were not seen with older releases. (Q01799513)
- VLACP operation has now been improved to properly bring down the port under some specific failure situations not previously covered, so that the traffic fails over to another link of the MLT/SMLT thereby preventing traffic loss. (Q01752976)
- On ERS 8600 the SMLT ID will no longer continue to be displayed under IP->RSMLT tab in JDM, if the corresponding MLT is deleted. (Q01802676)

IP Unicast

- With ECMP enabled, all packets having the same source and destination IP and ingressing on an R module port will now properly follow the same one route. (Q01745613)
- The default route now works properly even after multiple failovers when ERS8600 is operating in HA mode. (Q01744096)

BGP

- While configuring BGP Peers or Peer Groups if multiple words are used in a string, the configuration is no longer lost upon sourcing the configuration file or upon a switch reset. (Q01773745)
- BGP packets will now be sent with DSCP value of 0x30 and TOS value of 0xC0. (Q01726133)

Multicast



- If the ERS 8600 is connected to an external router on an extended VLAN, upon disabling/re-enabling PIM on ERS 8600 there will no longer be traffic loss. Also the assert messages from the ERS 8600 are sent out with correct metric and preference. (Q01690284-01)
- In a PIM SSM configuration there is no longer duplicate traffic when the DR switch resets and comes back, or if there is a local interface on the DR connected to the source which is disabled and enabled. (Q01783623)
- ERS 8600 will no longer be occasionally unstable when PIM is disabled on its IST Peer in an SMLT based configuration. (Q01760775)
- If the IST session should go down, multicast data packet coming in on the IST_MLT physical ports no longer are forwarded to the CPU, but are instead properly forward by hardware. (Q01772951)
- If all the ssm-channels are simultaneously deleted and resulting in many simultaneously leave messages being sent by the switch, the switch is no longer unstable. (Q01814078)
- ERS 8600 is no longer unstable when network load balancing mode is configured as multicast. (Q01790517)

JDM

- RMON history could previously not be displayed for certain configurations; this is now resolved by using JDM 6.0.9.0 or higher. (Q01755752)

8. Outstanding Issues – To be resolved in future release

- If DNS server is configured on ERS 8600 and a user does a cut/paste of commands onto a telnet session only the first command within the cut/paste will be executed. (Q01784202)
- The **pcmcia-stop** option is displayed twice in the command list in CLI. This is a minor display issue. (Q01785465)
- When modules are inserted in ERS 8600 chassis, BGP can lose its Peers for approximately 10 seconds. The situation is seen when the hold-timer is configured less than 5 seconds and keep-alive timer is configured less than 2 seconds. Please reference CSB 20080087818, under ERS 8600 (CS-LAN). (Q01825956)
- Advanced R-module filter operation may not work as expected after a system reboot, as the order of the filters may not function properly. This is only seen under some very specific filter combinations. If seen please contact Nortel support who can provide a temporary work-around to have the filters work as expected. (Q01814530)
- On ERS 8600 in a SMLT set-up when the static route next-hop is disabled and enabled, in some specific scenario there could be loss in connectivity. (Q01827994)
- If an IP prefix list is created via JDM or CLI with a name that contains blanks and it is used for configuring igmp access control entry, upon a re-boot the configuration will not load properly. For now, it is recommended not to create IP prefix list name with a blank in the name. (Q01829843)
- Protocol packets, such as OSPF hellos or VRRP keep-alives may not be sent out while saving the configuration file in a dual SSF/CPU system, when the backup SSF has a flash/HW problem and the save config to standby fails. (Q01827936)
- SNMP/JDM connectivity is lost after the hsecure flag is set to true (enabled) and the ERS 8600 is reset, because setting the hsecure flag to true, currently forces the block-snmpp flag to operate as being set true, regardless of the actual setting. Planned future operation is to allow block-snmpp flag to be set to either



state, true or false, independent of hsecure flag, as is currently allowed for other boot flags, such as telnet or tftp, or any other boot flags whose operation hsecure might affect. (Q01811342/Q01811343)

- Similar to the above documented CR, the ssh-pass-auth system parameter (config sys set ssh pass-auth <true|false>) can not be set true, when hsecure boot flag is enabled or set true. (Q01818955)
- Under some rare conditions, an LACP enabled link that is part of a LAG, but is then determined by LACP to be non-operational, can form a valid connection to the remote device outside of the LAG if that connection is based on link status only; this operation can lead to a loop between the original LAG, and the newly formed link connection. An example would potentially be if LACP for both the local and remote ports is disabled at the port level. For such an operation, it is recommended to first administratively disable one of the ports, before disabling LACP on both of the ports. (Q01819762)
- ERS 8600 can be unstable when garbage characters are hitting the console port at a high rate and at the same time the user tries to login to the switch. This situation was seen with a terminal server connection to a console port, and with improper terminal server configuration. At this time proper configuration of devices connected to console ports is recommended, but this operation will be enhanced by better protection in a future (4.1.7.0) release. (Q01825851)

9. Known Limitations – Operation not to be changed

- When one of the I/O modules is removed from the IST peer, the number of FDB entries on the master and slave (requires HA configuration) will not be same. The problem happens only when deletion and learning of FDB entries are happening simultaneously. (Q01788373-01)
- Users are recommended not to create more than 117 legacy destination filters having a destination IP addresss of 0.0.0.0. If a user create more than 117 of such filters it could result in the IST session going down and traffic loss. Users can look to use an alternative filter configuration, such as using source IP filters. (Q01787079)
- No event notification is logged in system log-file file when the pcmcia-stop command is executed, as system logging to the pcmcia has already stopped. Users using the pcmcia-stop command should be aware of this activity prior to, or when, using this command. (Q01774929)
- For PIM, a maximum of 100 candidate RPs should be configured. Configuring more could lead to improper PIM operation. (Q01797061)
- On R modules a mismatch in the duplex settings may cause unidirectional traffic flow. Users are warned against improper auto-negotiation settings, which can lead to this behavior. Reference CSB 2008008723. (Q01587376)
- In a PIM network having a VLAN extended over multiple switches there could be scenarios where the incoming port is also displayed in the outgoing port list. However multicast data traffic will not flow out of the incoming port and there is no functional impact. Also the pruned port may stay in prune pending state in some specific scenarios. (Q01752755)
- Upon enabling IP filters on nearly all the ports of the switch at the same time, the CPU utilization may reach 100% for couple of minutes and user will not be able to access CLI during this time. It is suggested not to perform such an operation, but instead to enable filters on port-by-port basis. (Q01776709)
- Using custom egress queues, i.e. a queue size greater than 8, traffic shaping can cause traffic to be dropped on all queues if this custom egress queue receives traffic for more than the configured value. This problem does not occur when the number of balanced queues is reduced to 10 or below or if the min and max rate values are changed from their default value of 0/0 for all balanced queues. Either of these can be used as a workaround, but if using custom queues, it is usual for the min and max rates for the



balanced queues to be adjusted from default values of 0/0 to start with, and is the suggested configuration. (Q01735632)

- The SDM FW does not support IP Multicast operation (PIM-SM), nor are there any current future plans for adding this functionality. (Q01345746)
- Under High CPU utilization the CP-limit feature may not always function properly. This is one reason Ext-CP Limit was introduced. Both CP-Limit and Ext-CP Limit can be enabled on any system at the same time; these are complimentary functions. (Q01804815)
- On ERS 8600 running in STP mode, the ports which were part of non-default STG are not added to default STG once the switch is reset in RSTP mode. A switch reset to RSTP mode, is best to be re-configured from default settings. (Q01810281)
- When a Passport 15K is connected to ERS 8600 using OC3, the Passport 15K does not detect the loss of signal when the port is disabled on the ERS 8600. (Q01472705)
- Do not mix a Non E-Series legacy 8608SX/GBIC module with an 8630GBR while creating an MLT. (Q01615420)
- Enabling of RSTP on ERS 8600 may cause periodic increase of CPU utilization as compared to normal STG. However the CPU utilization can be decreased by increasing the **hello-time** in RSTP configuration. (Q01800935)
- Under RSTP port path cost, a port changes its path cost from the default value of 200000000 to some value depending upon port speed, which is often determined via auto-negotiation. Now if the port is then disconnected, the port path cost will not return to the default value, but instead will remain at the last value associated with last speed. Now if the switch is re-booted, the port's path cost will now return to the default value of 200000000. (Q01840692)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

The following TCP ports are open in ERS 8600 by default

- Port 80 for Web server
- Port 21 for FTP
- Port 514 for Syslog
- Port 513 for rlogin
- Port 23 for telnet
- Port 111 for SUN RPC

The following UDP ports are open in ERS 8600

- Port 67 for BOOTP
- Port 514 for Syslog
- Port 162 for SNMP TRAP
- Port 161 for SNMP
- Port 69 for TFTP
- Port 111 for SUN RPC

These ports are open purposely. Although the state is shown as Listen for these ports, no connections will be accepted on these ports unless that particular daemon is enabled. Ex: Telnet on port 23. We cannot dynamically start and stop the daemons at runtime, so we need to keep them running right from the start. (Q01730700/Q01806655)



- The CLI command “config bootconfig sio modem 8databits true” sets the number of data bits per byte to 8, while a setting of false sets the number of data bits to 7. The documentation in the user manual Getting Started (P/N 313189-F, rev 00) states the opposite. False is the default parameter value. (Q01827281)
- The CLI command “rstp ethernet slot/port rstp protocol-migration” can be set to a value of true, but once set a ‘show info’ for this same command, will display false. This is proper and intended operation. For further details reference the 802.1w standard regarding ‘mcheck’. (Q01840688)
- The following documentation is corrected in “Network Design Guidelines - Ethernet Routing Switch 8600 Software - Release 4.1” (313197-E Rev 00) under section “CP-Limit considerations with SMLT IST”

The CP limit default settings are

```
default state= enabled (SMLT links)
default state= disabled (IST)
default multicast packets-per-second (pps) value=10,000
default broadcast pps value=10,000
```

(Q01819869)

APPENDIX

MIB changes in 4.1.6.0 release.

```
rcPortAdminSpeed OBJECT-TYPE
SYNTAX  INTEGER {
    none(0),
    mbps10(1),    -- 10Mb/s
    mbps100(2),  -- 100Mb/s
    mbps1000(3), -- 1000Mb/s
    mbps10000(4) -- 10Gb/s
}
MAX-ACCESS      read-write
STATUS           current
DESCRIPTION      "Indicate this port's speed."
DEFVAL          { mbps10 }
::= { rcPortEntry 14 }
```

Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch

8600

Software Release 4.1.5.4

1. Release Summary

Release Date: 19 Dec 2007

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

2. Important Notes before Upgrading to This Release

Release 4.1.5.4 is a replacement maintenance release for Release 4.1.5.0; Release 4.1.5.0 is no longer web posted. Reference CSB 2007008059 Rev1 on the Nortel support web site. The one change in Release 4.1.5.4 as compared to 4.1.5.0 is that support for ER 01660093, new Global VLACP parameters, has been removed. Support for this new functionality is now targeted for Release 4.1.6.0 (end of March, 2008 delivery). Other edits to this readme, versus the original 4.1.5.0 readme, are shown in bold italic lettering.

If upgrading to 4.1.5.4 from any other 4.1.x release besides 4.1.4.0 or any pre-4.1.x release (3.7.x or 3.5.x), then CR 1596219 still applies, such that the user must reconfigure their OSPF MD5 keys, and save the configuration afterwards. Reference this CR in the 4.1.4.0 Readme. If upgrading from 4.1.4.0, nothing must be done in regard to this, if the reconfigure and save steps were previously done.

CR 1617338 still applies, in that a system without a Mezz image present should not be boot with Mezz flag enabled. Any system without a Mezz image present, even if a Mezz is installed, should have the Mezz flag option disabled.

IMPORTANT NOTE for SMLT Users: 4.1.5.4 introduced CR 1745278 - smltRemote log message maybe sent every 10 seconds continuously for edge box base vlan MAC. This message will NOT affect system operation, but instead of being a log message (sent to both log file and SysLog Server) the message should have been a trace message, as it was meant for internal debug use only. This will be changed in the next release, but for 4.1.5.4 the message can be ignored, as again the message has no operational affect. To reduce the amount of messages sent to the log file and SysLog Server, suppress these SMLT messages with the msg-control feature, via the following configuration:

```
sys set msg-control control-interval 30
sys set msg-control max-msg-num 2
sys set msg-control enable
sys set msg-control force-msg add smlt
```

Resulting in this:

```
CPU5 [01/21/08 20:59:19] SW INFO msgControl: messages starting with 'smlt' suppressed.
```

For upgrades from pre-4.1.x release (3.7.x or 3.5.x) to 4.1.5.4, the LACP status should be checked (especially if used) as the global default setting for LACP was changed in 4.1.x based code from prior default setting of enabled to new default setting of disabled. Please also note in the new feature section, a change to the VlacpEnable global command.

NOTE: If you are a customer who loaded 4.1.5.0 software and enabled the new Global VLACP parameters, and you are moving to 4.1.5.4 code, you must first manually edit your configuration file to remove the lines below before you upgrade, or else your switch, with default settings, will load the default configuration not your actual running configuration. The two lines are:



```
#  
# VLACP CONFIGURATION  
#
```

```
vflacp time-out scale 3          -> delete this line  
vlacp mac address xx:xx:xx:xx:xx:xx -> delete this line
```

3. Platforms Supported

Ethernet Routing Switch 8600 modules in 8003, 8006, 8010, and 8010co chassis.

The following modules are not supported in the 8003 chassis:

- 8630GBR
- 8648GTR
- 8683XLR
- 8692SF/CPU
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are now supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4154.img	Boot monitor image	1073380
p80a4154.img	Runtime image	8761926
p80j4154.dld	Run-time image for R modules	1272680
p80c4154.img	3DES	55928
p80c4154.aes	AES (this image includes the DES image)	26112
p80a4154.mib	MIB	3328453



p80a4154.mib.zip	MIB (zip file)	533135
p80a4154.md5	md5 checksum file	1442
p80p4154.dld	8600 POS module image	701771
p80t4154.dld	8600 ATM module image	906024
p80m4154.img <i>Firmware images</i>	8600 image for the SuperMezz card	8862255
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4154.pkg	SSL cluster upgrade	5988896
p80s4154.img	SSL boot monitor	7528448
p80s4154.upgrade	SSL upgrade instructions	1481
p80s4154.install	SSL installation instructions	2895
p80s4154.diag	SSL diagnostics	19460381
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004
<i>SDM Firewall Images</i>		
NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image	
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO	
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package	
<i>SDM TPS Images</i>		
Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image	
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO	
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1	



Nortel_TPS_IS_Upgrade
_xxx_Upgrade-10.md5

IS upgrade download verification
file.

Nortel_TPS_DC_Upgrade
_4.5.0_to_4.5.1_Upgrade -47.sh

Upgrade script (patch) to upgrade
TPS DC from 4.5.0 to 4.5.1

Nortel_TPS_DC_Upgrade
_4.5.0_to_4.5.1_Upgrade -
47.sh.md5

DC upgrade download verification
file

5. Version of Previous Release

Software Version **4.1.4.0**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.6.0

This software release supports the Web Switching Module (WSM) running version WebOS 10.0.34.0

7. Changes in This Release

New Features in This Release

Global fdb-filter support for R modules (ER Q01528780)

This feature provides the functionality of global fdb-filter for R-modules using the existing CLI command, which previously was only supported with classic non-R modules. Previously for R-module cards we were achieving this behavior using the R-module advanced filter commands. Now the global fdb-filter will be applicable to all the I/O modules in the chassis.

The global fdb-filter command allows the user to selectively add a list of source MAC addresses which the user wants to discard. This feature also allows configuration of multicast MAC address. User will not be allowed to configure System MAC, static MAC, learned MAC, STP BPDU MAC, SONMP MAC and an all zero MAC address. Once a MAC entry is added into global discard list, it cannot be configured statistically or learnt on any VLAN/port. Alternatively, once a MAC is statically configured or learned, these MACs are not allowed as options for the global fdb-filter command. No bridging and routing will be performed on any packets from these MAC addresses on any VLAN. Wild card entry is not supported. The addition of MAC address to the global discard list is done via the following CLI command:

config fdb fdb-filter add <mac>

This command also allows the user to selectively remove the list of MAC addresses from the global discard list. Once any MAC Entry is removed from the global discard list, bridging and routing can now be performed on all packets from this MAC address on any VLAN. The deletion of particular MAC address from the global MAC discard list is done by the following CLI command:

config fdb fdb-filter remove <mac>

The list of MAC address added to the global MAC discard list is displayed via the following CLI command:

config fdb fdb-filter info

All the limitations and restrictions of legacy global FDB is applicable for this feature also. The addition of this functionality for R-modules, does not restrict 'normal' Advanced R-module filters in any way. The user will still be able to configure and use Advanced R-Module filters, be they In or Out Port filters or In or Out VLAN based filters, just as they did prior to the addition of the global fdb-filter functionality for R-modules.



There is existing JDM, CLI and MIB support for this functionality.

Detection of FAD/SWIP Errors (ER Q01505098/Q01728593)

This feature was introduced to enable software to detect FAD/SWIP problems based on the parity errors detected on the 869X SF. When this feature is enabled, and the count of parity errors exceeds the user set threshold, appropriate action is taken, such as logging an error message, notifying using traps and/or disabling the failed 869XSF (monitor mode). This can be configured on both master and slave CPU.

The following CLI commands have been introduced to support this feature.

config bootconfig parity-errors <enable/disable>

This command enables or disables parity error monitoring.

config bootconfig parity-errors set <Size>

This command will set the threshold of the number of parity error count to trigger action.

config bootconfig parity-errors action-869xSF-disable <true/false>

This command can be used to disable the CP card when excess parity errors are encountered.

JDM support is not provided for this feature in this release. The following MIBs have been added for this feature. The exact MIB definitions are provided in the Appendix section.

rcnTmuxParityError

Choice of CLIP to be advertised in topology (ER Q01654904)

In the previous software release, there was no option available to advertise the circuitless IP interface in the topology discovery packets. Topology discovery packet always uses the static management IP address. This feature when enabled allows an ERS8600 to advertise any one of the configured CLIP IP in the topology discovery packets.

ERS8600 now provides a configurable option to the user for enabling the provision of choice of topology-ip and also the instance of the CLIP which will be used. When configuring the choice of CLIP-instance mentioned, ERS8600 will first check the existence of the corresponding CLIP to be advertised in the topology. If this feature is not used, the ERS8600 continues to use the current behavior.

The following CLI commands are provided in order to enable this feature.

config sys set force-topology-ip-flag <true/false >

This command is used to enable or disable advertising the CLIP in the topology discovery packets. The default value is set to false (keeps current prior behavior).

config sys set clipid-topology-ip <clip-id>

This CLI command is used for providing the instance of CLIP to be advertised in topology discovery packets.

NOTE: The CLI syntax for this command is shown wrong under config sys set info. The current CLI syntax shows CirclessIpId which is wrong, and such syntax will not be recognized. One must use the syntax noted above to program some Clip-Id; this is shown properly under config set sys ?. This syntax will be corrected in Release 4.1.6.0 (March, 2008).

config sys set info



This CLI command can be used to display the status of the topology-ip-choice flag, as well as to show the instance of the CLIP mentioned for clip-id-topology-ip choice (shows under invalid parameter/entry name CirclessIpd).

JDM support is not provided for this feature in this release. The following MIBs have been added for this feature. The exact MIB definitions are provided in the Appendix section.

```
rcChasForceTopologyIpFlagEnable
rcChasCircuitlessIpd
```

VLACP Global Configuration (ER Q01660093) – Support for this functionality has been removed.

NOTE: JDM support for VlacpEnable is now re-enabled with Release 4.1.5.4. Changes made to this variable from JDM will continue to take effect, just as with prior code releases.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

- MIB value for ifSpeed is now displayed correctly for any 10 Gigabit interface.(Q01174158-01)
- Radius passwords up to 32 characters can now be configured or forwarded to a RADIUS server. (Q01587270)
- There will be an explicit message logged when a user makes an unsuccessful log-in attempt using SSH. (Q01638642)
- A user will no longer be able to delete the default SNMP-v3 username configuration. (Q01669001)
- The same custom banner created and displayed on the Master will now be displayed correctly on the Standby/Slave SSF/CPU. As well, disabling or enabling the default banner and adding a custom banner is no longer allowed on the Standby/Slave SSF/CPU. (Q01669927)
- Connecting to the Backup CPU via rlogin no longer causes system instability. (Q01690288)
- Inserting a corrupted PCMCIA into an ERS8600 no longer causes system instability. (Q01716404-02, Q01736713)

Platform

- During software upgrade or downgrade the 8683XLR module will no longer intermittently timeout and will load its configuration properly. (Q01677993)
- Packets received on any R-module having a destination MAC starting from A or 2 will no longer be sent out on the ingress port the packet arrived on (reflected packet). (Q01684395)
- Ports belonging to an 8648GTR and 8648TXE I/O modules can now be configured in the same MLT. User is required to set the same speed and interface type on these ports. Proper MLT configuration



from a consistent configuration of all ports in the MLT from a speed and duplex view is the responsibility of the user, as there is no software check to prevent mis-configuration (Q01670186)

- FPGA firmware upgrade of a R-module I/O card can now be upgraded successfully if the I/O card is present in Slot 10. Previous any FPGA firmware upgrade would fail for modules in slot 10. (Q01678422)
- The port link errors stats are now cleared properly when the counters are reset, via the command config sys set action resetcounters, or equivalent JDM/SNMP command. (Q01578488)
- Upon booting an ERS 8600, the PCMCIA card will now be functional after reboot even when there are a very large number of messages being transferred to the PCMCIA card (system logfile). (Q01785442)

Layer 2 switching

- MSTP traffic is now forwarded over an interface with forceportstate set to enable, upon a switch reboot. (Q01673927)
- Upon switch reboot ERS 8600 now properly forwards traffic over a MLT which has ports with spanning tree disabled. (Q01678420)
- Traps which are sent out for unknown MAC violation now include the MAC address of the offending device in all scenarios. (Q01644511)
- Under MSTP, when CIST is disabled all MST instances are now also disabled. A User will no longer be able to enable a MST instance if its associated CIST state is disabled. (Q01742111)
- When RSTP/MSTP is enabled for a switch, all spanning tree is now disabled on all Single Port SMLT (SLT) configurations, which is the proper STP setting for any SMLT edge port. (Q01741831)
- When a port with operational status down and STP disabled is added to a MLT in RSTP mode, packets will no longer loop over the MLT links. (Q01687705)
- When a static ARP entry associated with a Multicast MAC is created the packets associated with this IP address will now go out only on ports configured for the static ARP entry. (Q01645944)
- In association with NLB operation, when IP ARP multicast-mac-flooding is enabled and the number of ports where a multicast MAC packet should be sent out is limited via configuration of static-mcastmac addresses, the multicast MAC packet received on an IST is now forwarded on other ports. (Q01637863)
- Ports which are brought down due to cp-limit, link-flap, loop-detect or SLPP will now be re-enabled only after the auto-recovery-delay time expires. Auto recovery timer is a global timer and not specific to ports and hence there is a possibility that the ports are re-enabled by auto-recovery in less time than the globally configured auto-recovery-delay time. (Q01731886)
- Hardware MAC records on R-modules will now be relearned without an interface reset after the hardware record was deleted due to any normal process. This operation could have previously led to improper communication and excess latency. (Q01756726-01, Q01756753-01, Q01735232)
- SLPP PDUs will now be sent with a modified source MAC address. This change is in order for SLPP to function properly in all scaled environment. The new SLPP MAC will now be programmed to contain the source Vlan ID. The last 2 bytes of the Source MAC address will now contain the Vlan ID. This new operation alleviates the need to be concerned with SLPP network design as previously called out in CR Q01578936 (see 4.1.3.0 Readme). (Q01612186)



- Source and destination classic IP traffic filters with mask value of 1 to 7, which were previously not supported, will now no longer be allowed to be configured. (Q01628617)
- CP limit functionality will now bring down the ports properly when the port receives broadcast or multicast traffic at rate exceeding the configured values. Also 802.1p broadcast packets with priority 7 are no longer recognized as multicast packets. (Q01717469-02, Q01720807-02)

MLT/SMLT

- FDB entries learnt over a non-SMLT link are now properly purged on both of the IST Core peer switches when the non-SMLT link is brought down. (Q01669984)
- Static ARP entry added for a host connected to an IST peer switch will no longer be deleted when the peer switch loses connectivity to the host. (Q01716187)
- If there is an additional MLT in parallel to an IST MLT between a pair of Dual Core IST switches (unusual configuration), the MAC hardware records now get removed properly once the FDB entry ages out. (Q01667281)
- The designated port of an IST MLT is no longer displayed as "Null" on one of the aggregation switches. Designated port within an MLT is the port that would handle STP, which is not applicable to any IST MLT as STP is automatically disabled within such an MLT. (Q01671049)
- If the RSTP flag is enabled globally, when an SLT link is broken the SMLT status is no longer incorrectly displayed as SMLT on both aggregation switches. (Q01666744)
- In an SMLT environment with VLACP enabled on any 8608 type [E/M] I/O modules, VLACP now recovers properly upon HA failover. (Q01766647)
- In a scaled SMLT environment, IST no longer flaps as a result of losing SMLT Hello messages when a very large number of SMLT messages are being exchanged between the IST peers. (Q01776485-01)
- The CLI command "show ip rsmlt info" now will display the SLT id correctly when more than 8 SLT are assigned to a RSMLT enabled VLAN. (Q01675726)

IP Unicast

- ERS 8600 now allows configuration of IP address on a VLAN which has tagged Brouter port as a member. (Q01666040)
- ERS 8600 does not allow the configuration of Brouter IP on an untagged port which is already a member of an IP enabled VLAN. (Q01666045)
- ERS 8600 will now forward all traffic belonging to the same source/destination pair to the same next hop of an ECMP group. (Q01667206)
- IP traffic classic filters will now function properly after a change in port status (up/down/up) or a reboot; this situation was probably introduced in some 4.1.x code, as this previously work in all 3.7.x releases. (Q01640420)
- Addition and deletion of VLANs associated with an ACL are now managed properly, which previously could cause routing issue between the VLANs. (Q01751220)
- Adding and deleting routes continuously no longer causes system instability. (Q01734830)



- A few specific situations associated with VRRP GARP packets which could lead to ARP and FDB entries not being programmed with the correct port information, are now handled properly. (Q01645567)
- IST no longer flaps on receiving packet with an unknown destination address. (Q01762359-02)
- On ERS 8600 IPv6 ping now works properly for tagged packets also. (Q01640148)

OSPF

- ERS 8600 will now remove Type 5 LSA for cases which require Appendix E processing (RFC 2328), when the route source is removed. (Q01678901-01)
- On ERS 8600 with HA mode enabled, primary and secondary OSPF MD5 keys created on the Vlans will now be displayed properly even after multiple HA failovers. (Q01593252)

BGP

- If two BGP peers have static route for the same network they are learned as alternate routes on the peer, then these routes are now removed from alternate route if the static route goes down. (Q01731834)
- BGP peers will no longer continue to advertise the static route that has been deleted from both the BGP peers. (Q01725946)

Multicast

- The CLI command “show ip igmp snoop” will now display all the VLANs which have IGMP enabled. (Q01658847-03)
- In a multicast over SMLT set-up, when the RP is present on the Non-DR switch, and data is flowing to the Non-DR, the hardware records are now timely deleted once the sender stops sending data. The non-deletion of these records could potentially lead to a hardware resource issue. (Q01674098)
- ERS 8600 will now generate a trap when IP mroute resource exceeds the pre-configured values. Previously only a log message was properly generated. (Q01686594)
- IGMP sender's port information will now be updated properly when the sender is moved from one port to another. (Q01679389)

8. Outstanding Issues – To be resolved in future release

- Using custom egress queues, i.e. a queue size greater than 8, for traffic shaping can cause traffic to be dropped on all queues if this custom egress queue receives traffic more than the configured value. This problem does not occur when the number of balanced queues is reduced to 10 or below or if the min and max rate values are changed from their default value of 0/0 for all balanced queues. Either of these can be used as a workaround, but if using custom queues, it is usual for the min and max rates for the balanced queues to be adjusted from default values of 0/0 to start with. (Q01735632)
- Connecting CPU management port to an Ethernet I/O port and setting the speed and duplex setting to 100 Full on both sides causes the management port to come up as 100 Half after a switch reboot. Workaround would be to configure both management port and Ethernet I/O port with auto-neg enabled, versus disabled, and not fix the speed and duplex, which is not required. (Q01737868)
- Changing the bootp status to enable or disable from JDM causes management connectivity to be lost. However, the same problem is not seen from CLI. Please use CLI as a workaround. (Q01746388)



- In an SMLT environment receiving IPv6 traffic, the MAC address of a host connected to the edge box may incorrectly be programmed to point to the IST port instead of the SMLT port when the Neighbor goes into stale state. In such a scenario, IPv6 traffic may take long time to recover after failover. (Q01752114)
- Issuing a Control-X while in telnet prompt may reboot the switch. At this time, please do not issue Control-X within a telnet session. (Q01752502)
- When a multicast receiver resides on the same port of a sender, the same interface is shown in the list of incoming port, outgoing port and prune pending port. Showing as outgoing port is incorrect. (Q01752755)
- When using PCAP to capture packets, and a large file is transferred between two hosts, PCAP will only function once, and for the next attempt stats will display PCAP packets as being dropped at hardware. This is due to lack of free buffers on Slave/Standby SSF/CPU. (Q01754856)
- Trace Level off option is currently not present in the R-Module trace level commands. In order to turn off tracing user has to reset the level to zero for each of the modules which previous has trace enabled (Q01756714-01)
- Single Fiber Fault Detection does not work properly when an 8630GBR module is connected to an OM1200/OM1400. For any other configuration, VLACP should be used instead, if any SFFD issues are seen. (Q01757462)
- Non-DR IPv6 OSPF neighbor remains in incomplete state and never recovers when any one of its OSPF neighbors is reset. (Q01752119)
- When RMON is enabled to gather history statistics and the bucket size requested is greater than 375, the RMON HISTORY tab under 'Graph' in JDM becomes unresponsive after a certain time interval. (Q01755752)
- When using an MSS15k PVG switch with 8600 ATM module with an OC3 MDA, at certain link utilizations, the InDropPkts starts incrementing followed by F5-OAM mismatches finally resulting in the PVC's going down. This is due to delay in the lack of free buffers and the ERS8600 responding to F5 OAM requests. (Q01497294)
- The default route stops functioning after multiple failovers when ERS8600 is operating in HA mode. Workaround for this is to use more specific static routes than the default route, or else not create multiple failovers for any HA configuration. Alternatively the switch could be changed to work in non-HA mode. (Q01744096)
- VRRP and OSPF transitions are seen on ERS8600 running PIM-SM when the number of IGMP group members reaches a high value. Workaround for this issue would be to configure IGMP access control at the edge (IGMP Proxy enabled) to reduce the number of IGMP messages. Once the number of IGMP group members is reduced, OSPF and VRRP transitions are no longer seen. (Q01732757)
- SMLT/SLT configured on ERS8600 remains in SMLT state, when all the ports from the SMLT/SLT VLAN are removed. (Q01751343)
- ERS8600 depends on some of the IGMP-MIBs which are currently discontinued by RFC 2933. Workaround for this would be to modify the MIB file in text format and replace the occurrence of IGMP-MIB with IGMP-STD-MIB. (Q01726908)
- A MAC address for which an fdb-filter is configured, is not displayed in the Forwarding table; the MAC should be displayed with type 'mgmt'. There is no functional impact due to this. (Q01497977-03)
- The following additional error message is displayed while creating an IP subnet based Vlan with a non-zero instance in RSTP mode. For example if:



config vlan 10 create byipsubnet-mstprstp x.x.x.x/mask, the following additional error code along with:
Error : Only default instance supported in RSTP mode

Error: Error code not found 74526832 (the additional error message; this message can be ignored) (Q01775159)

- When classic IP traffic filters are created with redirect next-hop option set to a some external interface, the ERS8600 continues forward traffic to the changed next-hop external interface after the IP address of the next-hop device is changed or replaced. (Q01774902)
- For Advanced ACL (R-module filters) the action of drop for next-hop unreachable when associated with the re-direct action is not working – packets are still forwarded. (Q01597484)
- Radius authentication fails disallowing users from connecting to the switch when an ERS8600 is booted with a Mezzanine card. (Q01786612)
- The amount of IPv6 ICMP errors that is allowed to be sent over the specified interval mentioned by IPv6 icmp-error-interval is set to 50 by default. With 4.1.x code the user is not allowed to modify this value as there is neither SW or JDM support for this. Rel 5.0.0.0 will allow user modification to this value, via a new IPv6 icmp-error-quota parameter addition; this new parameter will also have JDM support. (Q01348336-01)
- Move or Rename operations done on a non-existent source file to an existing destination file causes the destination file to be deleted. Be careful when using file management commands with existing files. (Q01791386)
- DDI SFPs are currently not recognized by ERS 8600, and will show a status of “unsupported” if installed in an 8630GBR. Despite this, the SFPs should function fine. (Q01785224)

9. Known Limitations – Operation not to be changed

- During an upgrade from 4.0.x release to 4.1.x release the usable flash memory can be upgraded from 16MB/40MB to 64MB. After the flash has been formatted the first command executed related to any file operation generates Error messages on the console. However, the system recovers automatically and this does not have any impact on functionality. (Q01681686)
- When packets egress out of a mirror port which has diffserv enabled, the data portion of the mirrored packet maybe corrupted for some interval of time. This situation does not affect the user packets. As well, if diffserv is set to disabled on the mirror port, the issue is not seen, which is the recommended choice. (Q01744991)
- Multicast packets ingressing on pre-E modules with the multicast MAC set the same as that of the enabled protocols like RIP, OSPF, VRRP or PIM will cause the CPU utilization to go high. Workaround for this issue would be to not use the same MAC address which the well-known protocols use when the switch has any pre-E modules installed. (Q01718994)
- When an ERS8600 has learnt more than 2000 FDB and ARP entries and all the entries age out at the same time, some stale ARP entries may continue to exist in the ARP table. To clear the stale entries a clear ip arp for all ports or all VLANs command can be used. (Q01732608)
- The "unknown-mac-discard activation enable" command should be executed as the last command of all port security commands; otherwise the "unknown-mac-discard" function will be activated with default parameters. In this instance the ERS8600 port will discard all unknown mac packets, even if autolearn has been enabled after activation. (Q01733255, Q01768943)
- Static routes **CAN** be configured with VRRP IP address as the next hop IP address. There are no known limitations with this configuration. The Rev 1 Readme stated that this configuration was not



supported, which was wrong. The actual issue turned out to be a mis-configuration, which caused traffic to incorrectly hit the CPU, causing high CPU utilization, but no network issues. (Q01735135)

- When copying a file from one CPU to another, if the target file already exists user is not prompted if they would like to overwrite the file or not. This is however the same behavior when a file is copied to the remote TFTP server from the Switch and the remote CPU is equivalent to a remote TFTP server. Care should be taken when performing copies so that incorrectly overwriting an existing file does not occur. (Q01764845)
- When LACP is configured on SMLT MLT ports and on rebooting one of the aggregation switches, the remaining MLT ports on the edge switch can go through STP state transitions which lead to traffic loss for 30 secs. In order to prevent this it is advised to configure smlt-system-id globally, or else to correctly configure STP disabled at the edge switch. When LACP is used with any SMLT or RSMLT configuration the use of the smlt-system-id is now always recommended to be used as a best practice for best results. (Q01692648)
- When using ByIPSubnet based VLANs for use with G.729 voice traffic, packet loss has been seen. Choices are to either change the voice format to G.711 or not to use a ByIPSubnet based VLAN configuration and instead tag both the port and the user traffic. (Q01745558/Q01759995)
- When SMLT is configured with LACP, on booting one of the aggregate switches LACP re-converges on the other aggregate switch which might cause the Vlan interface to go down for a brief period of time. In an RSMLT environment, this will cause the RSMLT peer to go down and start the hold-down timer instead of backing up for the peer. In order to avoid RSMLT re-convergence in such a scenario, the smlt-system-id needs to be configured. When LACP is used with any SMLT or RSMLT configuration the use of the smlt-system-id is now always recommended to be used as a best practice for best results. (Q01764323)
- The designated port in a LAG (or MLT) is the one which runs STP. A change to the port status of this port will cause STP to be re-run, and with default values always, which will cause a traffic loss of 30 seconds. (Q01770793)
- ***For an 8630GBR if a 1000BaseT SFP is inserted, and auto-neg is set to disabled, the port will always show up (operational status up) even if there is no physical connection. Configuration of an 1000BaseT port with auto-neg disabled is not a supported configuration per the IEEE 802.3ab standard, but for the 8630GBR and the use of a 1000BaseT SFP we do not guardrail against this setting, therefore it is up to the user to not use such a configuration. (Q01755512)***
- The CLI command “config log logToPCMCIA <true/false>” was introduced to stop logging to PCMCIA gracefully only at runtime/boot. Hence modifications to this command will not be saved across a reboot. (Q01761564)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

- The behavior of BGP default route origination is as follows:

config ip bgp orig-def-route <enable|disable>

This CLI command enables the advertisement of default route to the neighbor(s) if the default route is present in the routing table.

config ip bgp neighbor x.x.x.x originate-def-route <enable|disable>



This CLI command is use to advertise the default route to the neighbor for which this flag is enabled. The default route need not be present in the routing table for this. This command should not be used if the global orig-def-route flag is enabled. (Q01521243, Q01588792-01)

When PIM is enabled on the SMLT IST peer, then on booting the a switch, the switch comes up with a default static route added with the next-hop as IST peer for first 60 seconds for faster IPMC forwarding. Therefore as now recommended in the best practice guides, the IST VLAN should have PIM-SM enabled for any SMLT/R-SMLT configuration that supports L3 IPMC with PIM-SM. (Q01694222)

For any mention of LX SFP specifications in any documentation, the correct launch power is -9.5 to -3.0 dBm.

- Enabling/Disabling tagging on a port re-enables STP on the port; this is the way the product has functioned since 3.0 code. (Q01771405)
- Layer 3 IPv6 is not supported on IST and [R]SMLT, only L2 IPv6 is supported.
- The following documentation is corrected in “Network Design Guidelines - Ethernet Routing Switch 8600 Software - Release 4.1” (313197-E Rev 00) under section “VRRP and other routing protocols” for “Figure 48 Sharing the same IP address”

In the figure, Router 1’s IP address should be 30.30.30.2 and Router 2’s IP 30.30.30.1 (Q01691202)

- A VLAN having both IPv4 and IPv6 address can be associated with more than one Ingress ACL and Egress ACL. (Q01775850)

MIB changes in 4.1.5.4 release.

```
rcnTmuxParityError NOTIFICATION-TYPE
  OBJECTS    {rc2kDeviceGlobalSlot}
  STATUS     current
  DESCRIPTION "A rcnTmuxParityError trap identifies a problem in the FAD/SWIP based on the number of
parity errors."
  ::= { rcTrapsMib 165 }
```

```
rcChasForceTopologyIpFlagEnable OBJECT-TYPE
  SYNTAX     TruthValue
  MAX-ACCESS read-create
  STATUS     current
  DESCRIPTION "Used to enable/disable flag which is used to
set the CLIP-ip as topology ip"
  DEFVAL    { false }
  ::= { rcChassis 53 }
```

```
rcChasCircuitlessIpId OBJECT-TYPE
  SYNTAX     INTEGER(1..32)
  MAX-ACCESS read-create
  STATUS     current
  DESCRIPTION "The Clip-Id set by the user to be used as Topology-ip"
  ::= { rcChassis 54 }
```

```
rcnMcInRecMoreThanThreshold NOTIFICATION-TYPE
  STATUS     current
  DESCRIPTION "An rcnMcInRecMoreThanThreshold signifies that
number of MC Ingress records has exceeded the
configured ingress-threshold."

  ::= { rcTrapsMib 163 }
```

```
rcnMcEgRecMoreThanThreshold NOTIFICATION-TYPE
  STATUS     current
  DESCRIPTION "An rcnMcEgRecMoreThanThreshold signifies that
number of MC Egress records has exceeded the
configured Egress-threshold."

  ::= { rcTrapsMib 164 }
```

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600 Software Release 4.1.4.0

1. Release Summary

Release Date: August 1st, 2007

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

2. Important Notes before Upgrading to This Release

Please reference the new version of the 4.1.x.x Upgrade Guide, P/N NN46205-400, for upgrade instructions.

Please reference the below chart for inter-product SMLT operational support guidelines for any 4.1.x.x release:

Topologies	8600
Triangle L2	YES
Triangle VRRP	YES
Triangle RSMLT	YES
Triangle Multicast	YES (IGMP at the edge)

Topologies	8600-8600	8600-5500	8600-8300	8600-1600
Square L2	YES	NO	NO	YES
Square L3-Unicast	YES	NO	NO	NO
Square RSMLT	YES	NO	NO	NO
Square IP-Multicast	YES	NO	NO	NO
Full Mesh	8600-8600 Only including PIM IP-Multicast			

3. Platforms Supported

Ethernet Routing Switch 8600 modules in 8003, 8006, 8010, and 8010co chassis.

The following modules are not supported in the 8003 chassis:

- 8630GBR
- 8648GTR
- 8683XLR
- 8692SF/CPU
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:



- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are now supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4140.img	Boot monitor image	1072078
p80a4140.img	Runtime image	8752430
p80j4140.dld	Run-time image for R modules	1271568
p80c4140.img	3DES	55928
p80c4140.aes	AES(this image includes the 3DES image)	26112
p80a4140.mib	MIB	3327018
p80a4140.mib.zip	MIB (zip file)	532817
p80a4140.md5	md5 checksum file	1489
p80p4140.dld	8600 POS module image	701771
p80t4140.dld	8600 ATM module image	906024
p80m4140.img <i>Firmware images</i>	8600 image for the SuperMezz card	8850961
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		



p80s4140.pkg	SSL cluster upgrade	5988896
p80s4140.img	SSL boot monitor	7528448
p80s4140.upgrade	SSL upgrade instructions	1481
p80s4140.install	SSL installation instructions	2895
p80s4140.diag	SSL diagnostics	19460381
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004
<i>SDM Firewall Images</i>		
NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image	
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO	
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package	
<i>SDM TPS Images</i>		
Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image	
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO	
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1	
Nortel_TPS_IS_Upgrade_xxx_Upgrade-10.md5	IS upgrade download verification file.	
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh	Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1	
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5	DC upgrade download verification file	

5. Version of Previous Release

Software Version **4.1.3.0**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.6.0

This software release supports the Web Switching Module (WSM) running version WebOS 10.0.34.0



7. Changes in This Release

New Features in This Release

RSMLT L2 edge support for Dual IST Core switch failure situation (ER 1277216)

RSMLT implementation does not use a Virtual IP address but instead uses Physical IP addresses for redundancy. At the same time, RSMLT can be deployed in either L3/routed configurations, or L2/edge configurations, where previously one might have used VRRP (and back-up master). Now previously, if there was a power outage or shutdown of both switches within a Dual Core IST pair and for some reason only one switch came back-up, clients using the powered-off switch's IP/MAC as their default gateway would lose connectivity to the network. In such a scenario, even through RSMLT is enabled on the switch it was unable to backup for the Peer as it was unaware of the Peer's IP/MAC address. With this new feature, the RSMLT Peer's IP and MAC addresses are stored in the config file and will be used upon next reboot, if the IST link does not become active and operational. Otherwise, the switches will learn from their Peer as normal; see below for additional information.

This feature can be enabled/disabled by the following CLI command:

```
config ip rsmlt rsmlt-edge-support <enable/disable>
```

When the configuration file is saved, if the rsmlt-edge-support flag is enabled and RSMLT peer is UP, the peer IP address and MAC address is also automatically saved.

The peer information is cleared by the following CLI command.

```
config ip rsmlt clear-rsmlt-peer [<vlanId>]
```

NOTE: if the peer information is cleared the switch could stop forwarding for the peer.

After both the Dual Core IST switches have come back-up, and if the IST comes up and is operational, if a RSMLT Peer enabled message is received from the peer, then normal RSMLT operation is followed. If the peer has either an IP or MAC change, then a new save config must be performed in order for the new information to be saved, and RSMLT L2 Edge support to operate correctly.

However, if the IST Peer up message is not received (for example RSMLT is not enabled properly) and the rsmlt-edge-forward flag is enabled, then first the RSMLT Hold down timer is started to allow routing protocols to converge; during this time user operation could be affected. Once the hold down timer expires, saved peer information is picked up and Switch starts to backup for the Peer by adding the previously saved MAC and ARP records. The Holdup timer is then started and once this timer expires the previously added MAC and ARP records are deleted and the switch stops backing up for the Peer, as the Peer is not running proper RSMLT for the VLAN. It should be noted that RSMLT is a per VLAN parameter, and therefore all affects are on a per VLAN basis, not necessarily per switch.

In L2 Edge support mode the Local values of the HoldDown timer (default value of 60 seconds) and HoldUp timer (default value of 180 seconds) will be used.

NOTE: this feature is supported only for IP RSMLT Vlan's and not for IPX RSMLT Vlan's.

SYSLOG Messages from CLIP or Management Virtual IP (ER 1528350)

A new command "ip-header-type", had been added under config sys syslog tree. This command can be used to set the IP header in the Syslog packets to default or circuitless-ip or management-virtual-ip and one of these choices will be used in the packet.

The choice of IP in the IP header of syslog packets can be set by the following CLI command:



```
config sys syslog ip-header-type <default|circuitless-ip|management-virtual-ip>
```

If the ip-header-type is set to default, then for syslog packets that are transmitted inBand via I/O ports the IP address of the VLAN will be used. For syslog packets that are transmitted out-of-Band via the management port the physical IP of the master is used in the IP header. If the ip-header-type is set to management-virtual-ip then for syslog packets that are transmitted out-of-Band only via the management port the virtual management IP address of the switch is used in IP header. If the ip-header-type is set to circuitless-ip then for all syslog messages (inBand or out-of-Band) the circuitless IP address is used in the IP header. If a user has configured multiple CLIPs, the first CLIP configured will be used.

The following CLI command displays the value of *ip-header-type* set.

```
config sys syslog info
```

Changes in SMLT Design (ER 1374217/ER 1422918)

This enhancement aims at improving the existing SMLT behavior so that better redundancy and resiliency is achieved in some specific fail-over [actually recovery] cases. Until now, when a Core IST switch boots up, the closet switch could detect a link up event before SMLT could converge. If the edge started to send data over the link before the convergence had completed this could lead to some amount of unnecessary traffic loss.

Now, while the switch is coming up all the SMLT ports are kept down until the SMLT protocol converges. The MAC and PHY are enabled once the SMLT protocol has converged. Proper forwarding records are now programmed in the I/O hardware before a closet switch starts sending traffic over the link, thereby ensuring sub-second to no traffic loss. This could cause a complete power cycle of both switches within a DUAL (or larger) IST Core to take a little longer to pass traffic, but previously any traffic sent from the edge was probably being lost anyway. In the case of a single switch power-cycle or switch re-boot within the Dual IST Core pair, the other switch is operational, and therefore the longer time to start forwarding has no user traffic impact.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

General

- The message of the day string now accepts up to 1024 characters which are inclusive of End of Line characters. (Q01614724)
- Any CLI session is now terminated upon link loss on both console and modem serial ports. This includes if the cable is removed or disconnected. The user would need to reconnect and re-login. (Q01544526-01)
- The value of Last Time Sent for RMON events is now displayed correctly via CLI. (Q01617831)

Platform

General

- CWDM SFPs are now correctly displayed via CLI or Device Manager. (Q01579140)
- IP Flow Information Exchange (IPFix) feature now works properly when configured on port 48 of slot 10. (Q01568603)



- Normal software informational log messages will now display the name of the task and some additional information regarding the task that has been killed. This applies mainly to SF/CPU switchovers in non-HA configurations, where killing specific tasks is normal operation. (Q01559759)
- Egress queue stats are no longer lost upon hot swapping of I/O modules. (Q01587379)
- ERS 8600 will no longer be reset when upgrading R-module with any firmware image, if the file name is missing from the specified path. Instead, if the filename is not present the error message "*No such file exist in flash or pcmcia /flash*" would be displayed

If an improper file name is given the following error message will be displayed
[04/11/07 00:31:31] COP-SW ERROR Code=0x10101010 Slot 3: foqUpdate Failed ERROR at or near XSVF command #1. See line #1 in the XSVF ASCII file"
(Q01583614)

- The Temperature alarm LED now changes color on Master and Slave SF/CPU at the same temperature (now 50 degrees Celsius for both). (Q01596103)
- The IST will no longer become unstable when the IP Flow Information Exchange (IPFix) collector becomes unreachable. (Q01541086)
- For classic module hardware port configuration registers could get corrupted in conjunction with SF/CPU abnormalities, and thereby cause traffic loss. There is now a software check to correct the situation should it occur. Previously a switch re-boot was required to correct the situation. (Q01585007)

Layer 2 switching

General

- Port(s) added to a SVLAN will not get removed from that VLAN after switch reset. (Q01656588)
- VLAN 4093 can now be configured only when a WSM module is not present in the chassis, and the configuration is now saved across a switch reboot. From prior configurations, VLAN 4093 must be deleted before a WSM module is inserted in the chassis, for the WSM to get properly initialized. (Q01590210)
- An IST which is configured in association with an R-module port previously could go down for few seconds when a gratuitous ARP for a multicast MAC was received. This was seen mainly in NLB environments, but is now resolved. However in this case the multicast MAC is learnt on IST port on one of the aggregation switches. Hence that switch will not be able to connect to the server. (Q01547987)
- When multiple Access Control List ACEs are simultaneously disabled from Device Manager, multiple COP software error messages are no longer seen on the console or in the log file. (Q01581660), (Q01581699)
- Multiple COP software error messages are no longer seen on the console or in the log file, in association with NLB configuration. (Q01589048)
- When multiple SrcMac based VLAN are created, the information about SrcMac VLAN with lower VlanId will now be displayed properly. However, the RSP memory limitation still exists. Please refer to known limitation section for more information. (Q01577517)
- When configuring VLACP, the ethertype value is only accepted in hex format. In the configuration file the ethertype values are now also saved in hex format. For pre-4.1.4.0 code releases the ethertype values were added in hex format but saved in the configuration file improperly in decimal format. This would then create an incorrect setting upon a switch reboot or re-sourcing of the current saved config file. If a user upgrades to 4.1.4.0 release with a configuration file of an older version of code, the ethertype values



may not be correct (if previously modified) and the user may have to reconfigure the value, in proper hex format. A save config will now, with 4.1.4.0 code, also save the value in proper hex format. (Q01487359)

IP Unicast

General

- In an RSMLT set-up when one of the Dual Core IST pair switches is reset, all non-local default static routes, whose next-hop is not reachable, now stays inactive even after a user does an “apply” to a route policy from Device Manager or CLI. (Q01521914)

OSPF

- When OSPF traps are enabled on ERS 8600, the trap messages will now contain the new improved 4.1.x based information about state transition between OSPF neighbors. This new information was previously contain in log messages, but was missing from trap messages. (Q01559787)
- OSPF MD5 authentication keys now are saved in a non-config type file based on VLAN ID rather than IfIndex. Prior OSPF MD5 keys will be lost upon an upgrade to 4.1.4.0. Users therefore must reconfigure their OSPF MD5 keys after an upgrade to 4.1.4.0, and then save the configuration. Please see the newer version of the Upgrade Guide for proper steps. (Q01596219)

Multicast

- The CLI command "show ip igmp snoop" now displays all the Vlans which have IGMP snoop enabled correctly. Also multicast records for the sender which has stopped sending multicast data will be removed properly. (Q01658847-01)
- On an ERS 8600 if the AllowMoreSpecificNonLocalRouteEnable parameter is set, and a more specific dynamic route is present which falls into a local route subnet, IP multicast data will now flow from sender to receiver. However, the receiver will receive traffic for more than two more-specific-non-local-routes in the routing table only if the network IDs of the more-specific-non-local-routes are different and the senders have non-overlapping IP. When the AllowMoreSpecificNonLocalRouteEnable parameter is used a warning message that portions of the network could become unreachable is generated. Use of this command should only be taken when complete knowledge of the possible affects are well understood. (Q01597677)
- On an ERS 8600 NLB packets configured for multicast mode [only] which are to be routed are no longer dropped by an R module port if the ingress and egress port is the same. (Q01565360)

8. Outstanding Issues – To be resolved in future release

- Radius passwords greater than 20 characters can not be configured or forwarded to RADIUS server. Future release will add support for 32 character password length. (Q01587270)
- MSTP traffic is not forwarded over an interface with forceportstate set to enable, upon a switch reboot. (Q01673927)
- With Device Manager version 6.0.5.0, or any SNMP Manager, a user can delete the default SNMP-v3 username configuration. This action will not be allowed in a future code release. (Q01669001)
- Ports belonging to an 8648GTR and an 8648TXE I/O modules can not be configured in the same MLT, even though the ports are set to same speed and interface type. (Q01670186)



- There is no explicit message logged when a user makes unsuccessful log-in attempt to an ERS 8600 using SSH. A log message will be added in a future release. (Q01638642)
- If RSTP flag is enabled globally, when an SLT link is broken the SMLT status is incorrectly displayed as SMLT on both aggregation switches. (Q01666744)
- The designated port of an IST MLT can be incorrectly displayed as “Null” on one of the aggregation switches. However this does not have any functional impact. The designated port in an MLT is the port that would run STP calculations, if required. (Q01671049)
- Traps which are sent out for unknown MAC violation do not include the MAC address of the offending device, when *violation-downport* and *violation-sendAuthenticationTrap* are enabled. (Q01644511)
- In association with NLB operation, when a static ARP entry associated with a Multicast MAC is created the packets associated with this IP address may go out on ports that are not configured for the ARP entry. This will create extra flooding, but should not affect the application being used, but this operation will be changed in a future release. (Q01645944)
- In association with NLB operation, when IP ARP multicast-mac-flooding is enabled and the number of ports where a multicast MAC packet should be sent out is limited via configuration of static-mcastmac addresses, the multicast MAC packet received on an IST is not forwarded on other ports. This can negatively affect NLB operation, and therefore this combination should not currently be used in NLB environments. (Q01637863)
- If the management port is disabled, its IP address is still used for topology advertisement. A configuration option for topology advertisements will be added in a future release. (Q01654904)
- If there is an additional MLT in parallel to an IST between a pair of Dual Core IST switches, the MAC hardware records may not get removed once the fdb entry ages out. (Q01667281)
- ERS 8600 allows configuration of Brouter IP on an untagged port which is already a member of an IP enabled VLAN. This configuration option should not be allowed. (Q01666045)
- ERS 8600 does not allow configuration of IP address on a VLAN which has tagged Brouter port as a member. This operation should be allowed. (Q01666040)
- The CLI command "show ip rsmlt info" does not display the SLT id correctly when more than 8 SLT are assigned to a RSMLT enabled VLAN. (Q01675726)
- ACE filters configured using JDM for egress queue assignments are lost after a switch reset. Current workaround is to configure via CLI, but do not specify the optional NNSC egress-queue parameter. (Q01673091)
- ERS8600 is not able to dynamically learn additional, or the removal of, NLB servers into a NLB server cluster, when nlb-multicast mode is set. To properly learn, the nlb-multicast mode must be disable and then re-enabled. This would be required if NLB servers are either added or deleted from the NLB server cluster. (Q01671933)
- The CLI command “show ip rsmlt info” incorrectly shows RSMLT as UP when all SMLT IDs associated with the RSMLT VLAN are deleted and switch is reset. (Q01651440)
- The CLI command “show ip rsmlt info” incorrectly shows RSMLT as UP even though no SMLT ID is associated to the RSMLT enabled VLAN. (Q01652193)
- The CLI command to display the hash calculation for the MLT currently displays an incorrect error message when the MLT has R module ports and VLAN associated with the MLT also has legacy ports as a member. Please note that prior reported CR 1581700 regarding potential improper operation of the



command was incorrect. It was that the source and destination TCP ports must be entered as decimal values, not hex. (Q01644166)

- User is able to create filters with an unsupported mask length of less than 8 bits. This situation is associated only with legacy cards. (Q01628617)
- ACL associated with a DSCP attribute does not work if any other attribute is configured along with it. (Q01691761)

9. Known Limitations – Operation not to be changed

- The CLI command “config ip ospf interface <ip address> disable” will disable OSPF globally. In order to disable OSPF on an interface please execute the following CLI command “config ip ospf interface <ipaddress> admin-status disable”. (Q01647814)
- In the case of Source Mac based VLAN associated with R-module ports the number of MAC addresses configured for a VLAN multiplied by the number of ports of an I/O module belonging to that VLAN can not exceed 75000. (Q01577517)
- On ERS 8600 when filters are created with copytoprimarycp enabled, the packets would be sent to primary CPU and CPU utilization would be high. When copytosecondarycp flag is enabled the packets are sent to secondary CPU but the CPU utilization will not be high. With both copytoprimarycp and copytosecondarycp enabled the packets will be sent to primary CPU but not to secondary CPU. The CPU utilization of only primary CPU would be high. Use of the copytoprimarycp option in any form is not recommended for any normal customer usage. (Q01541639)
- For radius accounting of SNMP, while sending the start/Interim/stop request, the packets will not contain the SNMP community string value. (Q01576692)
- If STG is enabled, there will be 30 second delay in forwarding for the whole MLT, when a new port is added or deleted from MLT. When a port is added or removed from an MLT, STG protocol is run over all the ports and it takes 30 seconds for Spanning Tree to put the ports in forwarding mode. (Q01674414)
- When SSH port is set to 65535, user is unable to make SSH connection to the switch. Do not use this port number for SSH connections. (Q01663948)
- MIB value for ifSpeed is not correct with 10 Gigabit interface. Currently, the MIB value 1410065408 represents a 10 Gigabit interface. (Q01174158-01)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

- The following documentation is added in “Upgrading to Ethernet Routing Switch 8600 Software Release 4.1” (Part Number -316674-C Rev00) under “reformatting the flash from 16 MB to 64 MB”:

After the format-flash command is issued, a reset command must be issued before the save file boot.cfg command or else erroneous error messages will be generated. This procedure only functioned properly via direct console access.



The proper steps for remote telnet operation are now outlined in the newly updated 4.1.x.x Upgrade Guide, P/N NN46205-400. (Q01611394)

- The following documentation is corrected in "Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) under "CR's fixed in this release 4.1.1.0" for CR Q01307380.

The "DirectedBroadcastEnable" feature is no longer enabled by default, and the 8600 does **NOT** forward Directed Broadcasts on a given IP Interface/VLAN by default. All new configurations contain the default value of "disabled." This change complies with RFC 2644.

- While configuring a static multicast MAC, the user has to create the static multicast MAC address before adding ports or an MLT to it or else add the ports or MLT along with the MAC address itself via the following CLI command:

```
Config vlan<vlanid> static-mcastmac add mac <value> [port <value>] [mlt <value>]
```

When the user adds an MAC address in the range 01:00:5e:xx:xx:xx a warning message is displayed, but the MAC address still gets properly added. (Q01567667)

- The CLI banner can be set to a customized value which can be as long as 1896 characters (inclusive of End of line chars). (Q01645319)
- The following documentation is corrected in Network Design Guidelines, Ethernet Routing Switch 8600 software Release 4.1" (Part number 313197-E) under "Hardware and supporting software compatibility"

The correct part number for 256MB CPU upgrade kit is DS1411016, not DS1404016 as listed on page 430. (Q01621944)

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch

8600

Software Release 4.1.3.0

1. Release Summary

Release Date: 13-Apr-2007

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

2. Important Notes before Upgrading to This Release

When upgrading any ERS8600 to 4.1.3.0 code, if a Mezz module is present, but the Mezz image is not, then the Mezz flag **MUST** be specifically set to false in boot.cfg or other boot file location (ex: pcmboot.cfg) . Please note that the default setting for the Mezz flag is enabled, and that previous 4.1.x releases did not have such a requirement. If the boot file is not configured in this manner, the system will look for the Mezz image and if the Mezz image is not present then the affects of CR Q01617266 will be seen. CR 1617266 is document in Section 8. (CR Q01617338)

3. Platforms Supported

Ethernet Routing Switch 8600 modules in 8003, 8006, 8010, and 8010co chassis.

The following modules are not supported in the 8003 chassis:

- 8630GBR
- 8648GTR
- 8683XLR
- 8692SF/CPU
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No. 316314-F Rev 01)

Note: R-series modules are now supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.



File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4130.img	Boot monitor image	1084319
p80a4130.img	Runtime image	8734661
p80j4130.dld	Run-time image for R modules	1269808
p80c4130.img	3DES	55928
p80c4130.aes	AES(this image includes the 3DES image)	26112
p80a4130.mib	MIB	3321305
p80a4130.mib.zip	MIB (zip file)	531934
p80a4130.md5	md5 checksum file	1489
p80p4130.dld	8600 POS module image	701771
p80t4130.dld	8600 ATM module image	906024
p80m4130.img <i>Firmware images</i>	8600 image for the SuperMezz card	8833826
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR (only applies to 8630GBR modules)	2284578
<i>SSL Images</i>		
p80s4130.pkg	SSL cluster upgrade	5988896
p80s4130.img	SSL boot monitor	7528448
p80s4130.upgrade	SSL upgrade instructions	1481
p80s4130.install	SSL installation instructions	2895
p80s4130.diag	SSL diagnostics	19460381
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004



SDM Firewall Images

NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package

SDM TPS Images

Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1
Nortel_TPS_IS_Upgrade_xxx_Upgrade-10.md5	IS upgrade download verification file.
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh	Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5	DC upgrade download verification file

5. Version of Previous Release

Software Version **4.1.2.0**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.3.0

This software release supports the Web Switching Module (WSM) running version WebOS 10.0.34.0

7. Changes in This Release

New Features in This Release

MSTP/RSTP Log Messages and SNMP Trap Implementation (ER Q01443993)

This enhancement implements the MSTP/RSTP traps and log messages for the various events triggered by MSTP and RSTP protocols to facilitate in trouble shooting these protocols.

For MSTP the log messages are added for following events:-

- Whenever a new VLAN or MSTI is created or the Region Name, Selector, or RevLevel is changed
- When a bridge is selected as the new CIST Root



c) CIST Port Role Selection

d) When a bridge is selected as the new CIST Regional Root

e) A bridge is selected as the new MSTI Root

f) MSTI Port Role Selection

g) When the port transitions to STP BPDUs from MSTP BPDUs or to MSTP BPDUs from STP BPDUs

h) When the port where edge-port was admin enabled, received a BPDU and reverted to OperEdge false operation

For MSTP the SNMP traps are generated for following events:-

a) When a bridge is selected as the new CIST Root.

b) When a bridge is selected as the new CIST Regional Root.

c) A bridge is selected as the new MSTI Root.

For RSTP log messages traps are added for following events:-

a) When a bridge is selected as the new Root Bridge

b) For Port Role Selection

c) When the port transitions to STP BPDUs from RSTP BPDUs or to RSTP BPDUs from STP BPDUs

d) The event occurs when the port where edge-port was admin enabled, received a BPDU and reverted to OperEdge false operation

For RSTP the SNMP traps are generated for following events:-

a) When a bridge is selected as the new Root Bridge.

Radius Support for SNMP (ER 1094361-01/CR 1094368)

After the introduction of SNMP-v3 implementation in 3.7 and subsequently 4.1, RADIUS authentication and accounting for SNMP required modifications. This enhancement aims at providing RADIUS accounting for SNMP.

The accounting will be done based on per SNMP Session which will record the duration of that particular session and the number of packets/octets received for this session. Accounting is done for every session. The user for any SNMP session has to be added as "snmp_user". At the beginning of any session, a start accounting message is sent to the RADIUS server. A stop accounting message is sent a period of time (based on the value configured for abort-session-timer) after the session is terminated. If the abort-session-timer is configured as 30 seconds (default value is 180 seconds) then a stop message is sent 30 seconds after the session is closed. The stop accounting message contains the duration for which the session was maintained and the number of packets/octets received for this session. If the session continues for a long period, then periodically (after every hour; non-configurable) an interim accounting message will be sent, containing the number of packets/octets received for that period for that session and the duration of the session.

Authentication is still done by the switch and not the Radius server. With the implementation of SNMP-v3, more powerful View based Access Control Model (VACM) is used to specifically permit or deny access to various OIDs. Since the security provided by the SNMP-v3 USM and VACM is quite powerful, radius authentication is not implemented for SNMP. Please note that SNMPv1 and SNMPv2 also use VACM for granting access to MIBS (OIDs) on our switch.



To enable radius accounting on the switch, a valid radius server must first be configured. To create a radius server, the following CLI command is used:

```
config/radius/server#
Sub-Context:
Current Context:
create <ipaddr> secret <value> [usedby <value>] [port <value>] [priority
<value>] [retry <value>] [timeout <value>] [enable <value>] [acct-port
<value>] [acct-enable <value>]
delete <ipaddr> usedby <value>
info
set <ipaddr> usedby <value> [secret <value>] [port <value>] [priority
<value>] [retry <value>] [timeout <value>] [enable <value>] [acct-port
<value>] [acct-enable <value>]
```

Radius accounting is enabled by using the following CLI command:

```
Under: /config/radius/snmp#
Sub-Context:
Current Context:
abort-session-timer <value> {30..65535}
acct-enable <true|false>
enable <true|false>
info
```

When the **enable** command is **true**:

If the **acct-enable** is **false**: accounting does not happen.

If the **acct-enable** is **true**: the on, off, start, stop, indata and outdata packets are seen on the radius.

When the **enable** command is **false**:

If the **acct-enable** is **made true**: a packet indicating that the acct-enable is turned on is seen on the radius. Accounting packets as mentioned above are not sent.

If the **acct-enable** is **made false**: a packet indicating that the acct-enable is turned off is seen, and as well, accounting packets are not seen.

For the accounting to happen, the user will have to make true both the *enable* as well as the *acct-enable* parameters.

Radius Accounting for SNMP-V1 & V2

Accounting is done based on the IP Address of the SNMP Client and the Community String used for accessing the switch. The following statistics are used for Accounting:

Acct-Status-Type: Indicates the status of the session whether it is alive or not.

NAS-IP-Address: Agent's IP address.

Acct-Session-Time: Duration of the session.

Acct-Input-Octets: Number of octets for the incoming SNMP packet for an authorized user.

Acct-Output-Octets: Number of octets for the outgoing SNMP packet for an authorized user

Acct-Input-Packets: Number of input SNMP packets for an Authorized user.

Acct-Output-Packets: Number of output SNMP packets for an Authorized user.



Client-IP-Address: IP address of the SNMP client to which the SNMP Manager is connected to.

Acct-Unique-Session-Id: Unique Id given for the existing session

The Input-packets/octets, & Output packets/octets are calculated only for authorized users, i.e only for the users to which switch has granted access. This data is calculated irrespective of the views associated to the community string, for e.g.: Input-packets is incremented of all authorized users irrespective of whether that user has access to a particular MIB or not.

Radius Accounting for SNMP-V3

Radius accounting for SNMP-v3 is same as V1 & V2 except that accounting is done based on per user name & IP address.

Please note the following which apply to both SNMP-v3 and SNMPv1/v2:

- a) Maximum number of Sessions that can run at a given time is limited to 100.
- b) If from a server multiple sessions are run with same community string then all those sessions will be treated as single session as they will have same IP/Community string.
- c) Accounting-On message will not be sent if the user has not configured a valid radius server before enabling radius accounting on the switch.
- d) Even after logging off from a session & re-logging in, previous session's data persists for some time based on abort-session-timer value.
- e) Traps are not accounted as a separate session cannot be created for traps.
- f) We do not support IPv6 Currently for RADIUS.
- g) When we open a session from JDM with SNMP v1/v2 login, two sessions are opened for the first time, but one of them is closed after N seconds, N again being the value configured for abort-session-timer, because Initially Both V1 & V2 packets are sent for authentication, then all the other info is sent are V2 packets. The session which was opened in the beginning for V1 is then closed.

MLT Hashing for IPv4 TCP/UDP traffic (ER 1497830-01)

The MLT Hashing algorithm has now been modified to be based upon Layer4 fields, for IPv4 UDP/TCP traffic for R-modules only, versus Source and Destination IP addresses. This hash change provides a better distribution of traffic for many environments, with no impact on R-module performance. MLT hashing for any module type is determined at the ingress port, although the affect is seen at the egress MLT. That is, the traffic ingressing a port determines the destination is an MLT and sends the traffic to the correctly hashed MLT port, regardless of MLT port type or even a mixed port type. Therefore in a mixed chassis, which algorithm that will be used will be based upon ingress port for the traffic, not the MLT configuration.

The new hash algorithm is as follows:

For IPv4 TCP/UDP traffic:

64-bit key = (SrcPort (16 bits), DstPORT (16 bits), DstIP (LSB 16 bits), SrcIP (LSB 16 bits))

For non TCP/UDP IPv4 traffic:

64-bit key = (DstIP (32 BITS), SrcIP (32 BITS))

The existing CLI command "config sys set hash-calc getmltindex" has been modified to allow src-port and dst-port as optional parameters. This parameter will provide the transmit or egress MLT index/port for R-module port ingress traffic.



Now, in some rare cases for TCP based traffic, the port displayed via the above CLI command and the actual port through which the traffic flows may be different. This is an outstanding issue (Q01581700) which will be resolved in a future release. Also the new hash algorithm is not applied to IPv6 traffic, just IPv4.

Note that the new hashing for IP traffic between a given source and destination IP address will be different for TCP/UDP packets and ICMP packets. Therefore the use of ping, in conjunction with the 8600 ping-snoop feature, is no longer always reliable to determine the hashed path taken by IP TCP/UDP traffic, if that hashing is performed by an R-module ingress port. If the hashing is performed by a legacy module, then ping-snoop functions just as before with other code releases.

Old Features Removed From This Release

None.

Problems Resolved in This Release

Switch management

General

- While configuring Day light saving time (DST) the date can now be entered in the Mm.n.d/hhmm or MMddhhmm format. However, DST will not work if start date and end date are configured in the same month or if the start and end date formats are different. (Q01531634)
- ERS 8600 no longer pads the OID of RFC 3418 traps with a value of zero. (Q01461504-01)
- In a system with and 8692SF/CPU with Mezz option installed and enabled, file system issue, which previously required a re-boot to clear, will no longer be seen. In the above type of configuration, rlogin boot flag should still be set to enable. (Q01475831)
- With Device Manger version 6.0.4.0, creation of a graph for filter statistics now works properly. (Q01529304)
- ERS 8600 now return proper value of MIB 2 structures for a 10GB card. However the value of ifspeed is incorrectly displayed as zero, as the counters for this interface are using 64 bit counters, while standard MIB 2 structures are based on 32 bit counters. There are no plans to change the ifspeed OID operation. (Q01540174)
- After a FTP session is initiated on an ERS 8600 switch, the RWA password is no longer saved in the cliilog.txt file. (Q01274075)
- In an HA configuration with the ERS8600 unable to load the dld image for R module cards, the system will no longer pick up an incorrect config file if the system also needs to load the backup-configuration-file. (Q01536179-01)

Platform

General

- ERS8600 Enterprise code now works properly with the 8692omSF/CPU. (Q01526505)
- RMON statistics for R modules are now displayed correctly. (Q01531702)
- The monitor port stats command will now show correct values of IN_UTIL and OUT_UTIL when the monitor interval is increased beyond the default value of 5 seconds. (Q01522756-01)
- CLI Command "clear port stats slot/port" now clears statistics for 10GE port. (Q01529652)
- On an ERS 8600 switch to load a Mezz image whose version is different from boot image version the Mezz image name has to be configured via the following CLI command



```
config bootconfig mezz-image image-name <Image-Name>
```

However in this case all the functionality might not work as expected. The following warning message would be generated upon boot time.

"Warning: The Mezz CPU was booted with Image Version: <version number>Incompatible with Boot Monitor Image Version <version number> .Not all functionalities might work as expected. Please either re-boot with proper Boot Monitor Image or load the proper Mezz/Runtime Image"

In this scenario for R-modules, the DLD version must to be the same as Mezz image version, or else the R-module will not come up. (Q01522508)

- 8630GBR ports will no longer flap or show errors, upon a disable of the port with certain SFP types in the connected device. (Q01498022)
- The Egress Queue Set for queue id 0 is now displayed properly in DM. (Q01490598-01)
- The CLI prompt will no longer be returned to the user until the actual FPGA update has completed, for either a direct console login session or a remote telnet/ssh/rlogin session. A spinning wheel has been added to inform the user that the FPGA update is still in progress. This new process will prevent the user from rebooting a switch too early in a FPGA upgrade situation, which previously might have made the R-module inoperable, and requiring an RMA. (Q01509339)
- The port statistics are now properly displayed for all 10GE modules. (Q01547182)

Layer 2 switching

General

- When both port state and VLACP are toggled on an SMLT edge port, the SMLT port now always comes back up properly and there are no connectivity issues, such as loss of OSPF adjacency across the IST MLT. (Q01525733)
- In a specific SMLT configuration where a combination of both SMLT and S-SMLT configurations are present, when a SMLT link is disabled there are no longer connectivity issues with the S-SMLT ports. A prior work-around to this situation was to configure the IST MLT with an MLT ID of 32, but that is no longer a requirement. (Q01538184)
- In an SMLT set-up with HA-CPU mode enabled and all IST connections running on R modules, when the SMLT Core Switch with the lower IP address associated with the IST VLAN is fully rebooted, the IST now come up properly. (Q01533253)
- If multiple SRC MAC based VLANs were configured the information of all the SRC MAC VLANs with lower VLAN IDs was not properly displayed. This situation is now resolved. (Q01534001)
- In a system which has multiple STGs, the STG port states are now displayed properly after an R-module is reset or upon an HA failover. (Q01479548-02)
- If an ACL with redirect-next-hop is configured and the specified next-hop is not available, after a ping to the unreachable next-hop all packets continue to be properly routed, assuming they have a valid route. (Q01492644)
- If an ACL with redirect-next-hop is configured on an ingress port of the switch and the specified next-hop is not available and global-action is set to mirror-count, the port no longer shuts down because of CP-Limit being invoked, and traffic will now flow as expected. (Q01493502)



- On an ERS 8600 switch which has dual CPU running in non HA mode, SMLT-on-SingleCP can now be enabled. (Q01512389)
- In a square RSMMLT set-up which has OSPF enabled, when the DR or BDR switch is reset the IST will come back-up properly. (Q01594982)

IP Unicast

General

- On an ERS 8600 unicast packets which are to be routed are no longer dropped by R modules if the ingress and egress port is the same; this requires the port to be tagged. (Q01511596)
- In an RSMMLT set-up when one of IST peer is reset, a non-local default static route whose next-hop is not reachable will now stay inactive. However the problem might be seen if the user does an “apply” to a route policy from Device manager or CLI. It is recommended that users do not create default static routes on both the OSPF neighbors pointing to each other, along with static to OSPF redistribute flag enabled, as this will create a potential route loop situation. (Q01521914/1522816)

OSPF

- In an SMLT set-up with HA mode enabled, when one of the other switches in the network is power cycled OSPF sync error messages are no longer seen on standby CPU of that switch. The same messages were never seen on master to start with. (Q01535376)
- ERS 8600 switch now handles OSPF Opaque LSA types 9, 10 and 11 properly. (Q01509489)

8. Outstanding Issues – To be resolved in future release

- On ERS 8600 IST can becomes unstable when IP Flow Information Exchange (IPFix) collector becomes unreachable. To not see this issue, disable IPFix globally or make sure some collector is reachable when using IPfix. (Q01541086)
- For an 8630 GBR, the CWDM SFP is incorrectly displayed by CLI or JDM as type equals LX, but the SFP operates correctly – this is a display issue only, with no operational affect. (Q01579140)
- On an ERS 8600 switch when a port is part of a subnet based VLAN, this port may drop traffic for this VLAN if filters are configured on any other port, as well traffic being of a frame size below 256 bytes and sent with a minimum inter-frame gap, meanings packets must ingress one directly behind the other. It is also believed this situation is only seen with certain E modules. (Q01557280)
- When OSPF traps are enabled the extra information about the state transition of adjacencies does not appear in the log messages. (Q01559787)
- When multiple Access Control List ACEs are simultaneously disabled from Device Manager, multiple COP software error messages are logged in the log file. Disable ACEs one at a time and the situation will not be seen. (Q01581660), (Q01581699)
- IP Flow Information Exchange (IPFix) feature may not work properly when configured on port 48 of slot 10. (Q01568603)
- On an ERS 8600 IST which is configured on an R-module port, the IST may go down for few seconds when gratuitous ARP is received which has two different MAC addresses associated to the same IP. This behavior has been seen in NLB set-ups with multiple NLB servers in same VLAN, and one NLB server set to unicast mode, while the other NLB server was being changed from unicast mode to igmp-multicast mode. It is suggest that if NLB operational settings are to be changed on a users server, that the servers be disconnected from the network (or the 8600 ports disabled) until the changes are made, and then reconnected. (Q01547987)



- The port link errors stats are not cleared when counters are reset. (Q01578488)
- On an ERS 8600 NLB packets configured for multicast mode [only] which are to be routed are dropped by an R module port if the ingress and egress port is the same. This requires the port to be tagged and connected to an L2 device. Moving the NLB device to either the same VLAN or different VLANs off of different ports will workaround this current limitation. (Q01565360)
- If you have a VLAN containing R-module ports configured to use VRRP Fast Advertisements, VRRP transition may happen and ports may go down on associated MLTs when the backup 8692SF/CPU is removed and then re-inserted in a dual 8692SF/CPU system, regardless of the HA-CPU setting. (Q01468173)
- In an HA configuration with the ERS8600 unable to load the dld image for R-modules, the system will pick up an incorrect config file if the system also needs to load the backup-configuration-file. (Q01536179-01)
- If a MAC address is added to the Source MAC based VLAN via SNMP error messages may be seen. Also information about the Source MAC based VLAN may not be correctly displayed by the CLI command "show vlan info src mac". (Q01577517)
- While upgrading a R-module with any firmware image, if the file name is missing from the specified path, it may cause the ERS 8600 to reset. Please make sure when updating R-module firmware images the proper and exact procedure is currently followed. (Q01583614)
- If a user configures a static multicast MAC address in the 01-00-5e-00-xx-xx range a proper warning message is generated, but the multicast MAC is not properly associated with the VLAN, but instead with just a port within the VLAN. (Q01567667)
- On an ERS 8600 with HA mode enabled, if a VLAN with OSPF has both secondary and primary MD5 keys, when the primary key is changed and HA failover is done twice both the keys are displayed as the primary key. (Q01593252)
- If a packet matches the port ACE and VLAN ACE which have the same mode and overlapping actions, the statistics of VLAN ACE are incorrectly incremented. (Q01592532)
- In a square RSMLT set-up which has OSPF enabled, when the DR or BDR switch is reset sometimes the IST may not come back-up properly. To resolve the situation toggle the IST status on the reset switch, to disable and then enable. (Q01594982)
- On a SMLT configured ERS 8600 the MD5 keys configured for OSPF authentication are not saved after first reset, if configured prior to IST configuration. The situation is not seen after the MD5 keys are added again and configuration file is saved. (Q01596219)
- On an ERS 8600 if the AllowMoreSpecificNonLocalRouteEnable parameter is set, and a more specific dynamic route is present which falls into a local route subnet, IP multicast data will not flow from sender to receiver. At this time it is recommended to not enable the above parameter, if any L3 IPMC routing protocol is also enabled. (Q01597677)
- For this release ERS 8600 pre-E/E/M module based chassis can only receive SLPP packets on 30 or less individual VLANs per unique SLPP sender. This limitation does not apply to any R-module only based receiving chassis. If this recommendation applies to your network design, then this affects the configurations on any chassis that is set for SLPP-tx enabled. Therefore, if you have an ERS 8600 pre-E/E/M module based chassis in your network, and you have SLPP enabled on more than 30 VLANs across your network, you must review your network design to ensure not more than 30 SLPP instances originate (are received) from the same source (chassis) even though some chassis (ERS 8600 pre-



E/E/M) may not even have SLPP enabled. If you only have R-Series modules in your network, you do not have to change any of your configurations or review your design. (Q01578936)

- When ERS8600 switch is booted with a Mezz module installed, but no Mezz image in flash or PCMCIA and the Mezz flag is set to true (by default Mezz flag is true) then the switch comes up with I/O modules working as expected, but the user can no longer ping/telnet/ftp/tftp to the switch through management port. If the Mezz image is available on flash or PCMCIA then there is no issue of ping/telnet/ftp/tftp to that switch. Also if the steps outline in Section 2 of this Readme document, reference CR 1617338, are followed the above situation will not be seen. (CR Q01617266)

9. Known Limitations – Operation not to be changed

- When an ACL redirect-next-hop filter is created with action unreachable command to permit, and if the next-hop is available but the ARP is not learnt on ERS 8600, packets would be forwarded, using a route in the route table, versus the configured next-hop. Workaround and recommendation for these types of configurations is to configure a static ARP entry for the next-hop. (Q01413767)
- On an ERS 8000 series switch a MAC is not added to the FDB table of a port based VLAN if the same MAC is added to a Source Mac VLAN. The problem happens only with Legacy cards. (Q01523549)
- Although STP port priority values are in the range between 0-255, only the following values are actually acceptable: 0,16,32,48,64,80,96,112,128,144,160,176,192,208,224,240. (Q01573415)
- While configuring IP Flow Information Exchange (IPFix) feature it is recommended to only use the v9 format selection. ERS 8600 sends out data using Netflow version 9 formats only. (Q01568558)
- When configuring VLACP the following standard ether-types should not be used. Also any well-known MAC address should not be used when configuring VLACP. There is no check in the code for this improper values, so care should be taken when configuring VLACP to use non-default values (CR 1583818):

Ethertype	Protocol
0x0000-0x05DC	IEEE 802.3 length
0x0800	IP, Internet Protocol
0x0806	ARP, Address Resolution Protocol
0x8035	DRARP, Dynamic RARP. RARP,Reverse Address Resolution Protocol
0x80F3	AARP,AppleTalk Address Resolution Protocol
0x8100	EAPS,Ethernet Automatic Protection Switching
0x8137	IPX, Internet Packet Exchange
0x814C	SNMP, Simple Network Management Protocol
0x86DD	IPv6, Internet Protocol version 6
0x880B	PPP, point-to-Point Protocol
0x880C	GSMP, General Switch Management Protocol
0x8847	MPLS, Multi-Protocol label Switching(unicast)
0x8848	MPLS, Multi-Protocol label Switching(multicast)
0x8863	PPPoE, PPP over Ethernet(Discovery Stage)
0x8864	PPPoE, PPP over Ethernet(PPP Session Stage)
0x88BB	LWAPP, Light Weight Access Point Protocol
0x8E88	EAPOL, EAP over LAN
0xFFFF	reserved



- On an ERS 8600 when an ACT has patterns for both ipv4 and ipv6, user will not be allowed to create ACL with pktType as ipv6. (Q01583818)
- If an ARP request is sent from an IP address with multicast Mac to ERS 8600, then it will not send out the ARP reply. The corresponding FDB entry will not be learnt resulting in loss of connectivity. (Q01595163)

Please also see Known Limitations Section of the Release Notes for Ethernet Routing Switch 8600 Software Release 4.1.1.0.

10. Documentation Corrections

While using any 8616 I/O blade with auto-negotiation enabled the re-autoneg flag must be set for proper interoperability. This flag can be set by the following CLI command (Q01563965):

```
config bootconfig flags 8616-reautoneg true
```

For the 8648GTR module or the 8648TX/TXE module, when we disable Auto-negotiation on a port and even without any link connected, the port will be set to a particular speed and the speed LED of the port glows in order to indicate the speed. The color of the Speed LED indicates the speed configured on the port. (Q01511924)

10/100/1000 speed cards: Auto-neg - Disabled, Link - Not Connected.

Speed set to 1000 - Speed LED glows Green.
Speed set to 100 - Speed LED glows Amber.
Speed set to 10 - Speed LED doesn't glow.

10/100 Speed cards: Auto-neg - Disabled, Link - Not Connected.

Speed set to 100 - Speed LED glows Green.
Speed set to 10 - Speed LED doesn't glow

- On an ERS 8600 switch with RSMLT enabled when a user configures a default static route with an exact preference value of 185, the static route is not saved in the configuration file. It is recommended not to configure default static route with preference value of 185 when RSMLT is enabled on the switch. (Q01587436)
- When auto-negotiation is enabled/disabled on any of the ports of the MLT, it will be enabled/disabled on all the ports of the MLT, and the MLT will go down and come back up. Since and IST is a special MLT, this operation also affects IST MLTs. Since SMLT is dependent on IST, when IST flaps, SMLT links will also flap. (Q01579329)

The following correction is done in document "configuring and managing security" (314724-E Rev-01) under section configuring SNMPv3. (Q01588855)

To configure SNMPv3

Step Action

There is a correction in the syntax of following commands:-

Create a group and assign the group to the user:
config snmp-v3 group-member create rdalton usm newgroup



Assign a mib view to the group.

```
config snmp-v3 group-access view newgroup pref usm authPriv newmibview write newmibview
```

Before executing step 7, please open Device Manager as mentioned in step 8.

- The following documentation is corrected in “Release notes for the Ethernet Routing Switch 8600 Software Release 4.1.1” (NN46205-402, 317177-E Rev 01) under section “CRs fixed in release 4.1.1” for CR Q01356928.

The issue with the remote link coming up before the 8630 GBR module is completely initialized has been resolved in this release. However for the fix to work user has to upgrade the FPGA firmware image on 8630GBR module to PIM=769.

- The mcast-mlt-distribution parameter and its operation is documented differently and inaccurately in several different manuals. The correct operation is listed below, which will be incorporated as part of the new re-write of the ERS8600 documentation set for the 5.0 Release in Dec., 2007. Please note that all MLT distribution, be it for unicast or multicast traffic, is a decision made by the ingress port, and therefore the port style make-up of the MLT itself has no affect on this operation. (Q01568460)

If the ingress port is a classic/legacy module port, and the parameter is enabled, distribution will always occur for IPMC traffic, be it L2 (same VLAN) or L3 (inter-Vlan).

If the ingress port is a R-module port and the egress MLT is in the same VLAN (L2 flow) distribution will occur only if IGMP snoop is enabled. For inter-VLAN or L3 flows, distribution will always occur for ingress IPMC traffic for a R-module port. In this configuration an IPMC L3 routing protocol must be enabled, or some static IPMC routing configured.

Distribution will never occur for a pure L2 multicast traffic (multicast destination MAC only), which has no L3 IPMC address, only if the ingress port is a R module port.

If IGMP snoop is enable for particular VLAN/Brouter port, CLI/JDM doesn't allow to configure L3 multicast protocol on that VLAN/Brouter port.

- A new version of the Installing an AC Power Supply in an Ethernet Routing Switch 8000 Series Chassis has been web post (P/N 312751-E, Rev 01), which contains the following change on various pages:

Nortel does not support 8001AC and 8002DC power supplies with R modules. R modules installed in the 8006, 8010, or 8010co chassis require the 8004AC, 8004DC, 8005AC, or 8005DC power supplies.

11. SNMPv3 Notification Configuration Guide (ER 1579640)

OVERVIEW

On Ethernet Routing Switch 8000 SNMPv3 Notification was introduced to replace traps. The older mechanism supported a simple configuration consisting of a trap receiver (typically an NMS), the SNMP version (either v1 or v2c) supported, and an associated community string. With Notification all the security that comes with SNMPv3 gets embedded into trap PDU's that NMSs expect to receive. Several ERS8600 features utilize the infrastructure Notification provides – traditional traps, SNMP inform messages, and RMON. With Notification comes two benefits – security and extensibility, but with a small price. So, this document is intended to address configuration intricacies since that's the cost of upgrading an ERS system that introduces Notification. The following configuration is not required if the user wants to continue with prior SNMPv1/v2 traps to some trap receiver.



NOTE: This document is taking a clean start and reconfiguring the security names from scratch, i.e. not assuming default security names. To perform such an SNMP-v3 reconfiguration from scratch and to start from a clean slate (all snmp-v3 tables empty) one would need to execute:

```
ERS-8606:5# config snmp-server bootstrap very-secure
WARNING: This command will remove SNMP community entries
WARNING: This command will destroy *all* existing SNMP configuration
Do you want to continue (y/n)?
```

If one chooses not to do the above (and hence keep the default entries in the snmp-v3 tables), the default usm security name "initial" should be deleted as it is a security hole if left in place, as anyone can connect with full rights without a password.

FUNDAMENTALS

Configuring Notification involves having knowledge of SNMPv3. While this document is not intended to cover SNMPv3 certain basic concepts will be touched upon.

- **(A) Notification Table**
 - Selects management targets which should receive notifications (including type)
- **(B) Target Address Table**
 - contains transport addresses to be used in the generation of SNMP messages
- **(C) Target Parameter Table**
 - manages all SNMP info as related to a target
- **(D) Notify Filter Profile Table**
 - associates a notification filter profile with a particular set of target parameters
- **(E) Notify Filter Table**
 - contains elements of a filter profile
- **(F) VACM Security-Group Table (rfc2575)**
 - maps a combination of securityModel & securityName into a groupName which is used to define an access control policy for a group of principals
- **(G) VACM Access Table (rfc2575)**
 - table of access rights for groups
- **(H) MIB View Table (rfc2575)**
 - locally held information about families of subtrees within MIB views

How to build a "Notification" configuration that determines what an ERS switch sends out

The following tables need to be configured: A, B, C, F, G, H. Here are the steps:

Configure an SNMP notification trap type (since this is part of the default configuration, this command does not need to be executed normally).

```
>> config snmp-v3 notify create Trap tag trapTag type trap
```

Define mib views as needed. In this example a default mib view is used (under normal conditions there's no need to execute this command).

```
>> config snmp-v3 mib-view create root 1 type include
```

Create group entries in the VACM. Associate mib views with each access entry.

```
>> config snmp-v3 group-access
>> create grp-1 "" snmpv1 noAuthNoPriv
>> view grp-1 "" snmpv1 noAuthNoPriv read root write root notify root
>> create grp-2 "" usm noAuthNoPriv
>> view grp-2 "" usm noAuthNoPriv read root write root notify root
```

Create users (one SNMPv1 & one SNMPv3)



```
>> create snmp-v3 community create indexForJohn John John
>> create snmp-v3 usm create Amy
```

Define security names and assign them to groups

```
>> config snmp-v3 group-member
>> create John snmpv1 grp-1
>> create Amy usm grp-2
```

Associate a target parameter entry with a security name (for mib access rights)

```
>> config snmp-v3 target-param
>> create parmList-1 mp-model snmpv1 sec-level noAuthNoPriv sec-name John
>> create parmList-2 mp-model usm sec-level noAuthNoPriv sec-name Amy
```

Associate target parameter info with target addresses. This calls out the snmp model, authentication/privacy (in the case of v3), and the type of notification to be sent.

```
>> config snmp-v3 target-addr
>> create NMS-1 10.10.10.10:162 parmList-1 taglist trapTag
>> create NMS-2 10.10.10.20:162 parmList-1 taglist trapTag
>> create NMS-3 10.10.10.30:162 parmList-2 taglist trapTag
```

Test Case #1:

Bring up an 8600. NMS stations NMS-1 & NMS-2 should receive a coldStart trap. The format should be SNMPv1 with a securityName of "John". NMS-3 should receive the same trap with a format of SNMPv3 and a securityName of "Amy".

How to build a "Notification" configuration that filters what an ERS switch sends out

The following tables need to be configured. B, C, D, E. Here are the steps:

Make sure target parameter entries exist (using the previous configuration).
Make sure target address entries exist (using the previous configuration).

Associate a notification based filter with an entry in the target address table

```
>> config snmp-v3 ntfy-profile create parmList-1 profile noLinkDown
```

Create notification based filters.

```
>> config snmp-v3 ntfy-filter
>> create noLinkDown 1 mask 0x80:00 type include
>> create noLinkDown 1.3.6.1.6.3.1.1.5.3 type exclude
>> create noLinkUp 1 mask 0x80:00 type include
>> create noLinkUp 1.3.6.1.6.3.1.1.5.4 type exclude
>> create onlySaveConfigFromRC 1 mask 0x80:00 type include
>> create onlySaveConfigFromRC 1.3.6.1.4.1.2272.1.21 type exclude
>> create onlySaveConfigFromRC 1.3.6.1.4.1.2272.1.21.0.37 type include
>> create onlySaveConfigFromRC 1.3.6.1.4.1.2272.1.21.0.61 type include
>> create onlyOspfIfStateChangeFrom1850 1 mask 0x80:00 type include
>> create onlyOspfIfStateChangeFrom1850 1.3.6.1.2.1.14.16 mask 0xff type exclude
>> create onlyOspfIfStateChangeFrom1850 1.3.6.1.2.1.14.16.2.16 type include
```

Special note when setting the mask for notification profiles:

The mask is optional. It can be set against every OID configured in the Notify FilterTable. The most significant bit of the mask is applied against the 1st dotted value in the OID. The subsequent bit in the mask is associated with the next dotted value in the OID (and so on up to a maximum of 16 bits). A bit value of 1 indicates an exact match to this value must occur; a bit value of 0 indicates a wildcard or any value in this field is accepted. Masks can be applied to both include and exclude "types".

Example: To filter on OID "1.3.6.1.4.1.*.1" a mask of 0xfd (1111 1101) must be applied. When the mask set contains fewer bits than the OID specified then the mask is extended with 1's (i.e. no wildcards are used). In the case where no



mask is entered an exact match of the OID specified is the filter (mask would be 0xff:ff). By default OIDs are considered "included".

Order in which traps are filtered:

1. The assigned "notify" value of user's mib-view (primary filter)
2. Only one profile per parameter list or user is configurable
3. The shortest (OID) rule in the notify filter table is executed first (the longest one is last)

Fine point: When setting filters be sure to pay attention to what's being done with the objects passed in trap OIDs. If the OIDs of the passed objects get excluded the entire trap will be excluded (In other words a trap sent cannot filter one passed object and not another).

Test Case#2:

Bring up an 8600 and force a data link in the system to go down. A link down trap should only be seen on NMS-3. Now, let's clear this condition by typing:

```
>> config snmp-v3 ntfy-profile delete parmList-1
```

Test Case #3:

Now, enter the following command. Then, force the same link back up again. You should see a link up trap on NMS-1 & NMS-2.

```
>> config snmp-v3 ntfy-profile create parmList-2 profile noLinkUp
```

Test Case #4:

Now, enter the following commands:

```
>> config snmp-v3 ntfy-profile
>> delete parmList-2
>> create parmList-1 profile onlySaveConfigFromRC
>> create parmList-2 profile onlyOspfIfStateChangeFrom1850
```

This configures all standard traps and only the save config trap from the PP8600 proprietary set of supported traps on NMS-1 & NMS-2. For NMS-3 all supported traps on PP8600 are configured with the exception of standard OSPF traps (other than the OSPF IF state change trap).

Now, attempt to save your configuration. You should see a save config trap on NMS-1 and NMS-2. If you attempt to generate a proprietary trap (say either a card up or down trap) you should see it on NMS-3 only. If you attempt to generate standard OSPF traps supported by PP8600 you should see all of them on NMS-1 and NMS-2. On NMS-3 you should only see the OSPF IF state change trap.

Additional Note: The SNMPv3 user configured here is not secure. Please review the document "Configuring and Managing Security" for more information regarding how to setup secure users.

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600 Software Release 4.1.2.0

1. Release Summary

Release Date: Feb-5th, 2007

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

2. Important Notes before Upgrading to This Release

Please note that in the original version of this Readme file, it was noted via CR 1501308 that the Daylight Savings Time (DST) configuration for the ERS8600 did not work – that information is incorrect. In fact DST configuration does function with the following exception:

DST will not function properly when DST-Start and DST-End are configured in the same month, except (again) if configured within the current month. Since DST-Start and DST-End are generally never configured in the same month, this operation will not be changed. Therefore there should be no current issues with DST configuration, with possible exception being related to CR 1531634 (see within this Readme), which will now be fixed as well in a future release.

CR 1501308 has now been removed from this Readme.

3. Platforms Supported

Ethernet Routing Switch 8600 modules in 8003, 8006, 8010, and 8010co chassis.

The following modules are not supported in the 8003 chassis:

- 8630GBR
- 8648GTR
- 8683XLR
- 8692SF/CPU
- 8683XZR

Please refer to the following documents for details on the Platforms Supported:

- Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)
- Installing Ethernet Routing Switch 8600 Switch Modules (Part No. 312749-K Rev 01)
- Managing Platform Operations Ethernet Routing Switch 8000 Series Software Release 4.1 (Part No. 315545-E Rev 00)
- Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis(Part No. 316314-F Rev 01)

Note: R-series modules are now supported in the 8010co chassis with a High Performance Backplane. Please see ReadMe for Ethernet Routing Switch 8600 Software Release 4.0.2.0.



4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category) for details on how to upgrade your Ethernet Routing Switch 8600.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4120.img	Boot monitor image	1081699
p80a4120.img	Runtime image	8725028
p80j4120.dld	Run-time image for R modules	1268612
p80c4120.img	3DES	55928
p80c4120.aes	AES(this image includes the 3DES image)	26112
p80a4120.mib	MIB	3319366
p80a4120.mib.zip	MIB (zip file)	531719
p80a4120.md5	md5 checksum file	1489
p80p4120.dld	8600 POS module image	701771
p80t4120.dld	8600 ATM module image	906024
p80m4120.img	8600 image for the SuperMezz card	8824185
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM8630GBR	2284578
<i>SSL Images</i>		
p80s4120.pkg	SSL cluster upgrade	5988896
p80s4120.img	SSL boot monitor	7528448
p80s4120.upgrade	SSL upgrade instructions	1481
p80s4120.install	SSL installation instructions	2895
p80s4120.diag	SSL diagnostics	19460381
<i>WSM Images</i>		



wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004
<i>SDM Firewall Images</i>		
NSF5100_2.3.3.0_SDM_R60.img	NSF Boot image	
NSF5100_2.3.3.0_SDM_R60.iso	NSF Boot ISO	
NSF5100_2.3.3.0_SDM_R60.pkg	NSF upgrade package	
<i>SDM TPS Images</i>		
Nortel_TPS_Intrusion_Sensor-SDM-v4.5.0-627-Install.iso	TPS Intrusion Sensor Boot image	
Nortel_TPS_Defense_Center-2x70-v4.5.0-627-Install.iso	TPS Defense Center Boot ISO	
Nortel_TPS_IS_Upgrade_4.5.0_to_4.5.1_Upgrade-47.sh	Upgrade script (patch) to upgrade TPS IS from 4.5.0 to 4.5.1	
Nortel_TPS_IS_Upgrade_xxx_Upgrade-10.md5	IS upgrade download verification file.	
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh	Upgrade script (patch) to upgrade TPS DC from 4.5.0 to 4.5.1	
Nortel_TPS_DC_Upgrade_4.5.0_to_4.5.1_Upgrade -47.sh.md5	DC upgrade download verification file	

5. Version of Previous Release

Software Version 4.1.1.1

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.3.0

This software release supports the Web Switching Module (WSM) running version WebOS 10.0.34.0

7. Changes in This Release

New Features in This Release

None

Old Features Removed From This Release

None.

AVAYA

Problems Resolved in This Release

Switch management

General

- When a save config command is executed, the saved config file is equal to the last runtime config used. Now, if the primary configuration file in the bootconfig (/flash/boot.cfg) is different from last runtime configuration file then the following warning message would be displayed on console, written to the log file, but not displayed on telnet/ssh logged in sessions. This is a warning to the user that should the switch boot, the config file that the switch will boot with, maybe different then the current running and saved config file.

"SW WARNING Choice Primary Node Config file is <file name>"

Where <file name> is the name of the primary configuration file. (Q01401503)

- The PCAP global configuration parameter AutoSaveFileName no longer allows a zero character file name to be configured. (Q01381856-02)
- Users will now be able to login to an ERS8600 using SSH with Public Key Authentication.(Q01483948, Q01483488)
- The following RMON traps, last system change and Ethernet octet stats, are now properly recognized by Device Manager. (Q01142068)
- The System Configuration parameter "#! flags system-monitor false" will now be displayed as TRUE when issuing the command "show config". Previously the flag displayed a value of false, even though it was operationally true. This parameter is not user settable as well. (Q01485344)
- If repeated Telnet or FTP sessions are opened against any 8600 In Band and/or out of band IP addresses with quick closures, over the course of time this could lead to system instability. This situation is no longer present with this code release. (Q01462275)
- Incompatibilities between Device Manager Version 6.0.3 and WSM modules running version 10.0.34.0 are no longer seen with this code release. (Q01491448)
- RADIUS authentication will no longer fail when users repeatedly log out before the RADIUS server sends the reply. (Q01406191)
- RLOGIN sessions which are kept idle are now logged out after expiry of time out period. (Q01410171)

Platform

General

- If the SuperMezz option is enabled and the switch is reset, ATM modules will no longer get disabled. (Q01463279)
- When multiple WSM cards are installed in the same chassis, upon a switch reboot all WSM cards will now be initialized properly. (Q01473364-01)
- The total available power for the 8005AC RoHS compliant power supplies, DS1405012-E5, P/N 321411, will now use a value of 1372W, which is correct for the supply running at 220 VAC, versus the prior usage of the default value of 690W. This makes the 8005AC RoHS power supply function like the non-RoHS version. This change will alleviate users from potentially seeing any erroneous warning messages regarding insufficient power for the system. At the same time this functionality always assumes the 8005AC, both the RoHS and non-RoHS versions, is running at 220 VAC. For systems that are run at 110 VAC, the calculation will be incorrect and this field should be ignored. As well the calculation for Power



Usage is not 100% correct as well. Neither of the later two situations has been addressed in this code release. (Q01495801)

- The below warning message is displayed when configuration changes are being made to the egress queue set:

WARNING: The egress-queue-set QoS change made will take effect only after the configuration is saved and the chassis is rebooted.

(Q01450460-01)

- The "show port info state" command will now reflect the time and date of the last state change for a port. (Q01474469-01)
- An ERS8600 with HA-CPU (hot standby) mode enabled, and containing an installed 8683XZR module, will no longer have any 8683XZR ports become inactive after an HA failover has completed. (Q01484430-01)
- 8630GBR when connected to SUN NIC with auto negotiation enabled no longer causes a connectivity issue to the end station with the SUN NIC. (Q01434934-01)
- For a dual SF/CPU non-HA mode (warm standby) configuration, with an 8683XZR card present, the system will no longer have the 8683XZR module become inactive when the master SF/CPU is reset. (Q01487003-01)
- The warning message displayed when the FPGA update command is executed has been modified as below to indicate to the user to wait for "FPGA UPDATE SUCCESS" message before executing any other command. In a future release of code, the CLI prompt will no longer be returned, until the actual FPGA update has completed.

WARNING: Starting the FPGA update process. DO NOT reset the card or box during this process. Please wait for FPGA UPDATE SUCCESS message.

(NOTE: Command prompt will be returned but update process can take up to several minutes)

(Q01501599-01)

- RoHS compliant GBICs are no longer displayed as GbicOther in CLI and Device Manager. (Q01453100-03)
- 8630GBR and 8683XZR cards will no longer become inactive when the master SF/CPU is failed on a switch with HA-CPU (hot standby) mode enabled. (Q01493870-01, Q01525078, Q01521770)
- An ERS8600 containing an installed 8683XZR module will now show proper configuration options for clock source and framing options for its WAN ports under all conditions. (Q01494672-01, Q01493665-01)
- For an 8683XZR module with a 10GBase LR/LW XFP running in WAN mode, the port will no longer go down after the port state is toggled. (Q01512208-01)
- On an ERS8600 switch having dual SF/CPU with a large configuration file and CLI log enabled, there are no longer any issues related to SF/CPU failover. (Q01525314)

Layer 2 switching

General

- Null Strings and Strings starting with a blank or a white space are no longer allowed to be configured as MLT names. (Q01478418-02)
- For Core IST switches in any SMLT configuration, ARP entries will now correctly point to SMLT ports versus the IST ports. (Q01476117)



- The configured STP state of a port in RSTP mode is now saved correctly in the configuration file. (Q01422525)

IP Unicast

General

- OSPF routes are now summarized properly when a Virtual link is present. (Q01444180-01)
- When a dhcp-fwd entry is created with the same IP as the VRRP IP, changing the VRRP configuration will no longer cause the CPU utilization to rise to 100%. (Q01473196-01)
- For an ERS8600 running OSPF and having an OSPF route with 32 bit mask that overlaps another OSPF route, the switch will no longer become unstable. (Q01491868)
- ERS 8600 will now send a trap and display a warning message when any VRRP IP address is assigned to any other device (duplicate IP) in the network. Previously duplicate IP warnings were only generated for physical switch IP address; this now adds warning support for VRRP IP addresses. (Q01352422-01)
- The value of configuration parameter **more-specific-non-local-route** will no longer be lost after a switch reset. (Q01491654)
- For an ERS 8600 with the **tagged-frame-discard** parameter enabled on a port, IPv4/IPv6 bridging on any protocol based VLANs associated with the port now works properly. (Q01427096)
- Dynamically learned routes can now be deleted from CLI via the following command "**config ip route delete <ipaddr/mask> next-hop <value>**" or from Device Manager via IP->IP-> Routes->delete. If an OSPF route is deleted the route will not reappear in the database unless there is network state change or if an SPF is run. (Q01475773)
- On an ERS 8600 switch executing the CLI command "**show log file tail**" followed by "**Q**" will no longer restart OSPF. (Q01506294)
- When one switch, in an IST Core configuration, is power cycled there will no longer be the potential for the IST to not form properly, especially in configurations where there are multiple OSPF adjacencies formed between the 2 IST Core switches. Additionally, there would no longer be unicast and multicast traffic loss. (Q01535355, Q01505335)

BGP

- Modifying the eBGP route preference to the recommended value of 15 no longer causes connectivity issues. (Q01407546)

IP Multicast

- ERS 8600 no longer allows PIM activity-chk-interval parameter to be modified. The following error message will be displayed when this operation is performed:

"Invalid Activity Check interval value".

Only the default value of 210 (seconds) is now allowed.
(Q01464950-01)

- On ERS 8600 with IP multicast enabled console messages such as:

servicePimAssertVifEntry: src 0.0.0.0grp 239.255.255.250

will no longer be erroneously displayed. (Q01469731)

- When (S, G) RPT entry is present in the DR and a receiver joins the group, the multicast routing table will now switch to the shortest path. (Q01464913-02)
- On ERS 8600 tagged multicast traffic that ingress and egresses the same R-module port will now be properly forwarded through R modules. (Q01478516)
- Multicast traffic is now forwarded to the downstream port when diff-serv is enabled and access-diffserv is set to TRUE on the ingress R module port. (Q01479569-01)

8. Outstanding Issues – To be resolved in future release

- The monitor port stats command will not show correct values of IN_UTIL and OUT_UTIL when the monitor interval is increased beyond the default value of 5 seconds. Currently do not change the monitor interval for proper operation. (Q01522756-01)
- On an ERS 8000 series switch a MAC is not added to the FDB table of a port based VLAN if the same MAC is added to a Source Mac VLAN. (Q01523549)
- Loopback testing commands are not working for R modules. Future support for the 8648GTR and 8683XLR and XZR is planned. Support on 8630GBR is not planned due to internal hardware restrictions. (Q01498160)
- On an ERS 8600 unicast packets which are to be routed are dropped by R modules if the ingress and egress ports are same. (Q01511596)
- ICMP packets with TTL of 1 are sent to the CPU as high priority packets, which under heavy usage could cause other high priority packets to be lost. This can lead to system instability. This is improper handling of these packets, which will be corrected in a future version of code. (Q01510545)
- CLI Command “clear port stats slot/port” does not clear statistics for any 10GE port. (Q01529652)
- When VLACP is toggled on an SMLT edge port, the SMLT port may not come back properly, which can lead to traffic passing issues, such as OSPF adjacency loss (RSMLT configuration). To workaround this, either do not toggle VLACP state, or not use VLACP on SMLT links. (Q01525733)
- In an RSMLT set-up when one of IST peer is reset, a non-local default static route whose next-hop is not reachable will becomes active. Do not currently configure non-local static routes in RSMLT configurations. (Q01521914)
- With Device Manger version 6.0.3.0, creation of graph for filter statistics does not work. This is a DM limitation and requires a change in a future version of DM for resolution. (Q01529304)
- The Egress Queue Set for queue id 0 is not displayed in DM. This is a switch code limitation, which will be resolved in a future release of ERS8600 code. (Q01490598-01)
- With Device Manager Version 6.0.2.0 and 6.0.3.0 (6.0.1.0 worked properly) the settings for the **duplex** parameter on a port cannot be changed. This is a DM limitation to be resolved in a future release. (Q01531475)
- RMON statistics for R modules are not displayed correctly.(Q01531702)



- ERS 8600 switch does not currently handle OSPF Opaque LSA types 9, 10 and 11 properly. (Q01509489)
- On an ERS 8600 switch the SNMPv3 group entries are not retained on switch reset if the group name includes space character. Do not create SNMPv3 group entry names that include a space. (Q01539098)
- ERS8600 Enterprise code will not currently work with the 8692omSF. (Q01526505)
- During the 60 seconds lockout period on master for a failed telnet, ssh, or rlogin session, users are also unable to make new telnet, ssh and rlogin connections to the slave, during the same 60 second lockout period. (Q01540836)
- On an ERS 8600 switch FTP, RLOGIN, TFTP are not supported for inbound IPv6 address. However these will be supported in future development release. (Q01492490,Q01476296,Q01474185)
- In an SMLT set-up with HA-CPU mode enabled and all IST connections running on R modules, when SMLT aggregation switch with the lower IP address associated with the IST VLAN is fully rebooted, the IST may not come up. Workaround is to disable and then enable the IST on the switch which has the higher IP address associated with the IST VLAN. (Q01533253)
- In an SMLT set-up with HA mode enabled, when one of the switch is power cycled OSPF sync error messages are seen on standby CPU of other switches. Workaround is to currently disable HA-CPU mode. (Q01535376)
- In an SMLT configuration where both SMLT and SLT are present, in a specific scenario when SMLT link is disabled the SMLT status on the Peer IST stays as "SMLT"; this can create connectivity issues. Workaround is to configure the IST MLT using an MLT ID with a higher value than any other configured MLT IDs, of which using an MLT ID value of 32 will always provide such a workaround. (Q01538184)

9. Known Limitations – Operation not to be changed

- On an ERS 8600 switch with R modules, the Ctrl-Rec-optimization feature is not supported, and the feature is recommended not to be used in a chassis with any R-modules present. (Q01292928-01)
- Single Fiber Fault Detection is supported on the ERS 8600 but should a user run into any issues with its operation, it is recommended that the user uses VLACP to achieve the same functionality. (Q01464922),(Q01347146-01)
- On an ERS 8600 switch with mezz flag enabled, it takes longer for the files to be copied from /PCMICA to /FLASH. (Q01526038)
- On the ERS8600 switch, the -f option cannot be used with SSH non interactive commands. (Q01393349-01)
- Using Device Manger version 6.0.3.0, a user is not able to do an **apply**, when multiple ACE's with mode permit are configured for an ingress ACL and the value is changed for either of the flags "copyToSecondaryCp" or "copyToPrimaryCp" for multiple ACEs. User will have to do apply for each ACE filter separately after changing the flag value. The copyToPrimaryCp flag is not recommended for general use as well. (Q01540698)
- Duplicate IP Multicast traffic over the IST can be seen in specific SMLT configurations. To avoid this situation, please reference the information contain in the Network Design Guide for Release 4.1 (Part No. 313197-E, Rev 00) page 305 – 309. (Q01539663, Q01270581)



Documentation Corrections

- In Release notes for the Ethernet Routing Switch 8600 Release 4.1.0 under “Hot-swapping the SF/CPU module or I/O modules”

When hot-swapping the active CPU/SF module in an Ethernet Routing Switch 8600 with dual SF/CPU modules (HA mode enabled or disabled), wait until the previous standby SF/CPU module is stabilized before inserting any other modules; this includes a replacement SF/CPU or reinserting the originally removed SF/CPU. This stability can be noted by the SF/CPU module displaying a login prompt on the console screen. If no console connection is available, wait for at least thirty seconds or until the previous standby CPU becomes master, whichever is longer before inserting the replacement SF/CPU module or before reinserting the originally removed SF/CPU module. (Q01489847-01)

- The default IP address for the management port is 192.168.168.168/24 for slot 3 of a 3-slot chassis and slot 5 of either an 8006, 8010, or 8010co chassis, and 192.168.168.169/24 for slot 6 of either an 8006, 8010, or 8010co chassis. The use of IP address 0.0.0.0 and subnet mask as 0 for management IP address is not allowed; therefore the management IP addresses can never be set to no IP address. If the 192.168.168.0 subnet is required for use for the I/O ports, then the management port IP addresses must be changed first. Once the IP addresses is changed, to change a second time, it maybe required to set the IP address back to the default values, and changes are then allowed to any address. (Q01492184)
- Configuring QoS and IP Filtering for Ethernet Routing Switch 8600 R Modules(Part No 318637-B Rev 00) under traffic shaping statistics, page 89, the following information regarding packets and memory pages should be noted and will be added in the next revision of the manual:

Each memory page is 512 bytes in length except the first page which is 144 bytes in length. The number of memory pages used for egress queuing will depend upon the packet size. A packet of size between 64-148 bytes requires 1 memory page, packet size between 149-632 bytes requires 2 memory pages, packet size between 633- 1120 bytes requires 3 memory pages and packet size between 1121-1518 bytes requires 4 memory pages.
(Q01539128)

- Configuring QoS and IP Filtering for Ethernet Routing Switch 8600 R Modules(Part No 318637-B Rev 00) under Updating Action parameters for an ACE following documentation needs to be added:

The optional parameter [mlt-index <index>] is not supported for multicast traffic, but only unicast traffic.
(Q01541824)

- Starting with JDM 603, there are two runtime generated property files: dm.xml and dm_devices.xml. They are placed in a hidden folder called "jdm" under user's profile. The first file holds the last opened device list which has the same as dm.ini with different format. The second one holds device property sets, one set per device. JDM uses the specific property set while communicating with the device. For example, SNMP status polling interval and timeout value can be set to different values for each device. It is a new feature in JDM.

This new feature keeps property files for each JDM user, not for each JDM installation. For example, on Window 2000, the files are located on d:\Profile\

The procedure to convert the dm.ini generated in JDM 59x to the property files used in new JDM 60x is as follows:

- locate the user profile directory, say d:\Profile\- go to jdm properties directory, cd d:\Profile\- if dm.xml and dm_devices.xml exist rename them or delete them



- copy the dm.ini from the previous jdm installation directory to this directory
- launch JDM and exit JDM
- the new dm.xml and dm_devices.xml are created by JDM and placed in the property directory say d:\Profile\\jdm
- launch JDM again, the saved device list is shown on the Open Last cascade menu.
- dm.ini is no longer in use for JDM as long as dm.xml exists and it can be removed

Because the property files are per user base, not per JDM installation, the conversion needs to be done only once, and the same files will be kept and reused for the subsequent JDM release (Q01546524)

- In some extremely rare cases, you will see the following error message "CPU [5|6] [[date] [time]] HW INFO <module type> card on slot <slot number> bootup timeout. Related configuration will be lost". This error message displays after a reboot, the side effect is that the ports of this module use the default configuration. It situation can happen if the line card comes up but is offline during the duration when the configuration is sourced. This situation causes the configuration to be lost on the ports of that card. If you experience such conditions, please source the current configuration file to make the module operational. (CR Q01395294)
- While configuring Day light saving time(DST) please enter the day in the < Mm.n.d/hhmm> format which is the n-th occurrence of day d in month m, along with hour and minute; only this format will work. (Q01531634)

11. SNMP Access policies - Clarifications

Since the Release 3.7.9, the behavior of Access Policies for SNMP has changed. This section will clarify the actions required during an upgrade from a previous release – 4.0.x to any of the 4.1.x releases.

How to configure access policy for snmp in 4.0.x?

1). Enable the access-policy globally:
config sys access-policy enable true

2). Create the policy:
config sys access-policy policy 2 create

The following parameters will be having default values when an access policy is created. User can change them if required. Below is the explanation of the parameters with examples:

- name -- Used to set the name of the access-policy
 - *config sys access-policy policy 2 name **policyName***
- policy enable – Used to enable or disable the policy
 - *config sys access-policy policy 2 **enable***
- mode – Used to determine whether access be allowed or denied to an incoming request if it matches the policy
 - *config sys access-policy policy 2 mode **allow***
- Precedence – If more than one policies are matched for an incoming request, the value of precedence determines which one of them would be applied. Lower the precedence value, higher the priority.
 - *config sys access-policy policy 2 precedence 10*
- network – If configured for a policy, then the policy gets applied only in case the subnet of the source ip address of the incoming access request matches the configured network address.
 - *config sys access-policy policy 2 network 198.202.188.0/24*



- host - If configured for a policy, then the policy gets applied only in case the source ip address of the incoming access request matches the configured host ip address.
 - `config sys access-policy policy 2 host 198.202.188.174`
- username – Used only in case of rlogin access.
 - `config sys access-policy policy 2 nortel`
- accesslevel – Determines the access level which the user should have if he is to be granted access. If the user's access level is greater than the one configured for the policy, he would be granted access or not depending upon whether access-strict is set to true or false.
 - `config sys access-policy policy 2 accesslevel rw`
- access-strict – Determines whether users with access levels greater than the one configured for the policy should be allowed access or not.
 - `config sys access-policy policy 2 access-strict true`

NOTE: With the above configuration for accesslevel and access-strict, only users with rw access would be allowed access. Had the access-strict been set to false, both rw and rwa would have been given access.

3). Login to the switch with public & private

Note: snmp service is default enabled in 4.0.x

How to configure access policy for snmp in 4.1.x?

1). Enable the access-policy globally:

`config sys access-policy enable true`

2). Create the policy:

`config sys access-policy policy 2 create`

The following parameters will be having default values when an access policy is created. User can change them if required. Below is the explanation of the parameters with examples

- Name -- Used to set the name of the access-policy
 - `config sys access-policy policy 2 name policyName`
- policy enable – Used to enable or disable the policy
 - `config sys access-policy policy 2 enable`
- mode – Used to determine whether access be allowed or denied to an incoming request if it matches the policy
 - `config sys access-policy policy 2 mode allow`
- Precedence – If more than one policies are matched for an incoming request, the value of precedence determines which one of them would be applied. Lower the precedence value, higher the priority.
 - `config sys access-policy policy 2 precedence 10`
- network – If configured for a policy ,then the policy gets applied only in case the subnet of the source ip address of the incoming access request matches the configured network address.
 - `config sys access-policy policy 2 network 198.202.188.0/24`
- host - If configured for a policy ,then the policy gets applied only in case the source ip address of the incoming access request matches the configured host ip address.
 - `config sys access-policy policy 2 host 198.202.188.174`
- username – Used only in case of rlogin access.
 - `config sys access-policy policy 2 nortel`



- `accesslevel` – Determines the access level which the user should have if he is to be granted access. If the user's access level is greater than the one configured for the policy, he would be granted access or not depending upon whether `access-strict` is set to true or false. This parameter is not applicable to snmp service in rel4.1.x
 - `config sys access-policy policy 2 accesslevel rw`
- `access-strict` – Determines whether users with access levels greater than the one configured for the policy should be allowed access or not. This parameter is not applicable to SNMP service in rel4.1.x
 - `config sys access-policy policy 2 access-strict true`

NOTE: Moving with the configuration example given above, only users with `rw` access would be allowed access. Had the `access-strict` been set to false, both `rw` and `rwa` would have been given access.

3). Enable `snmpv3` service in policy 2:

```
config sys access-policy policy 2 service snmpv3 enable
```

4). Add the `snmp-groups` and the security model to the access policy. The default `snmp-groups` and the security models for allowing access to private and public community strings are:

```
readgrp snmpv1
readgrp snmpv2c
v1v2group snmpv1
v1v2group snmpv2c
```

```
config sys access-policy policy 2 snmp-group-add readgrp snmpv1
config sys access-policy policy 2 snmp-group-add readgrp snmpv2c
config sys access-policy policy 2 snmp-group-add v1v2grp snmpv1
config sys access-policy policy 2 snmp-group-add v1v2grp snmpv2c
```

5). Login to the switch with public & private

6). If a new community name is to be granted access through a policy, then the `snmp-group` corresponding to the community string and the security model should be added to the access-policy.

```
config snmp-v3 create third nortel readview
config sys access-policy policy 2 snmp-group-add readgrp snmpv1
config sys access-policy policy 2 snmp-group-add readgrp snmpv2c
```

Upgrading from previous loads:

Upgrading from 4.0.x to 4.1.x

In 4.0.x the default community strings `public`, `private` and `secret` are mapped to access levels `ro`, `rw` and `rwa` respectively. The access policy check (if `snmp` service is enabled) is done against the configured policy's allowed access levels (access level & `access-strict` values).

In 4.1.x after the implementation of the `snmpv3` service option under the access policies, the check for the `snmp` service is on the basis of the `snmp-groups` along with the security levels allowed by the policy. The configured access level of the policy is not applicable to `snmp` service.

(NOTE: these access levels would still hold the same significance as before for services other than `snmp` like `ftp`, `telnet`, `rlogin` etc.). Each community string in 4.1.x would have a security name corresponding to it. This security name would be associated to an `snmp` group name under the `snmp-v3` group-member table, for different security models. Access to a particular community string is granted based on the group name and the security model. Please note that in 4.1.x, the `snmp` service in access policy is disabled by default.

For the upgrade from 4.0.x to 4.1.x to happen smoothly, the user is expected to do a '**save config verbose**' in 4.0.x before upgrading. During the upgrade from 4.0.x to 4.1.x, the `snmp-groups` get added to an access-policy



only in case the policy has snmp service enabled (which is the case by default). If the service is disabled, the snmpv3 service gets disabled and no groups get added to the policy.

For default communities public and private in 4.0.x with access-policy configured, their access state would be the same after upgrading to 4.1.x also. But for community secret, access would not be granted after upgradation. This is because community secret comes up as a default entry in 4.0.x, but not in 4.1.x. In 4.1.x there are just two default community entries created - public and private.

In 4.0.x the communities' public and private have access levels ro and rw by default. And in 4.1.x they have the security names readview and initialview respectively. Whenever a user upgrades from 4.0.x to 4.1.x, the command 'config sys access-policy policy id service snmp enable' would get executed through the config file (saved in rel4.0.x in verbose mode) for the policy. Now this command would enable the snmpv3 service for the policy in rel4.1.x and also add a set of snmp-groups with security models to the access policy. The snmp-groups would be added by default during the upgrade based on the access level to group mapping given below:

Access level	Access-strict	Snmpp-group ,Security model
RO	True	readgrp ,snmpv1 readgrp ,snmpv2c
RO	False	readgrp ,snmpv1 readgrp ,snmpv2c v1v2grp ,snmpv1 v1v2grp ,snmpv2c
RW	True & False	v1v2grp ,snmpv1 v1v2grp ,snmpv2c
RWA	True & False	-

TABLE 1

Now in 4.0.x, access is granted based on the community string to access level mapping. After upgrading to 4.1.x, access would be granted if the group names and the security levels corresponding to the community string are added to the policy. For the case where the access policy has rwa as its configured access level, we don't add any snmp-groups on upgrade because with this access level, only the community secret was being granted access in 4.0.x which is not available in 4.1.x by default.

EXAMPLES: Following are screen captures of some example configurations before and after upgrading from 4.0.x to 4.1.x

EXAMPLE 1): Upgrade with access level configured as ro, access-strict as false and snmp enabled in 4.0.x

In 4.0.x(before upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/2# info

Sub-Context: service

Current Context:

create :
delete : N/A
name : policy2



```
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readOnly
access-strict : false
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/2/service# info

Sub-Context:

Current Context:

```
http : disable
rlogin : disable
snmp : enable
telnet : disable
ssh : disable
tftp : disable
ftp : disable
```

In 4.1.x (after saving config in verbose mode and upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/2# info

Sub-Context: service

Current Context:

```
create :
delete : N/A
name : policy2
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readOnly
access-strict : false
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/2/service# info

Sub-Context:

Current Context:

```
http : disable
```



```
rlogin : disable  
telnet : disable  
ssh : disable  
snmpv3 : enable  
tftp : disable  
ftp : disable
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/2# snmp-group-info

snmpv3-groups :

<i>Group Name</i>	<i>Snmp-Model</i>
<i>readgrp</i>	<i>snmpv1</i>
<i>readgrp</i>	<i>snmpv2c</i>
<i>v1v2grp</i>	<i>snmpv1</i>
<i>v1v2grp</i>	<i>snmpv2c</i>

EXAMPLE 2): Upgrade with access level configured as ro, access-strict as true and snmp enabled in 4.0.x

In 4.0.x:(before upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/3# info

Sub-Context: service

Current Context:

```
create :  
delete : N/A  
name : policy3  
policy enable : true  
mode : allow  
precedence : 10  
network : 0.0.0.0/0.0.0.0  
host : 0.0.0.0  
username : none  
accesslevel : readOnly  
access-strict : true
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/3/service# info

Sub-Context:

Current Context:

```
http : disable  
rlogin : disable  
snmp : enable  
telnet : disable  
ssh : disable
```

In 4.1.x (after saving config in verbose mode and upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/3# info

Sub-Context: service

Current Context:

create :
delete : N/A
name : policy3
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readOnly
access-strict : true

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/3/service# info

Sub-Context:

Current Context:

http : disable
rlogin : disable
telnet : disable
ssh : disable
snmpv3 : enable
tftp : disable
ftp : disable

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/3# snmp-group-info

snmpv3-groups :

<i>Group Name</i>	<i>Snmp-Model</i>
<i>readgrp</i>	<i>snmpv1</i>
<i>readgrp</i>	<i>snmpv2c</i>

EXAMPLE 3): Upgrade with access level configured as rw, access-strict as true and snmp enabled in 4.0.x

In 4.0.x: (before upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/4# info



Sub-Context: service

Current Context:

```
create :
delete : N/A
name : policy4
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readWrite
access-strict : true
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/4/service# info

Sub-Context:

Current Context:

```
http : disable
rlogin : disable
snmp : enable
telnet : disable
ssh : disable
tftp : disable
ftp : disable
```

In 4.1.x (after saving config in verbose mode and upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/4# info

Sub-Context: service

Current Context:

```
create :
delete : N/A
name : policy4
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readWrite
access-strict : true
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/4/service# info



Sub-Context:

Current Context:

```
http : disable
rlogin : disable
telnet : disable
ssh : disable
snmpv3 : enable
tftp : disable
ftp : disable
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/4# snmp-group-info

snmpv3-groups :

Group Name	Snmp-Model
v1v2grp	snmpv1
v1v2grp	snmpv2c

EXAMPLE 4): Upgrade with access level configured as rwa, access-strict as true and snmp enabled in 4.0.x

In 4.0.x :(before upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/5# info

Sub-Context: service

Current Context:

```
create :
delete : N/A
name : policy5
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readWriteAll
access-strict : true
```

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/5/service# info

Sub-Context:

Current Context:

```
http : disable
rlogin : disable
```



snmp : enable
telnet : disable
ssh : disable
tftp : disable
ftp : disable

In 4.1.x (after saving config in verbose mode and upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/5# info

Sub-Context: service

Current Context:

create :
delete : N/A
name : policy5
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readWriteAll
access-strict : true

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/5/service# info

Sub-Context:

Current Context:

http : disable
rlogin : disable
telnet : disable
ssh : disable
snmpv3 : enable
tftp : disable
ftp : disable

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/5# snmp-group-info

snmpv3-groups :

Group Name Snmp-Model

EXAMPLE 5): Upgrade with access level configured as ro, access-strict as false and snmp disabled in 4.0.x

In 4.0.x: (before upgrading)



Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/6# info

Sub-Context: service

Current Context:

create :
delete : N/A
name : policy6
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0
username : none
accesslevel : readOnly
access-strict : false

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/6/service# info

Sub-Context:

Current Context:

http : disable
rlogin : disable
snmp : disable
telnet : disable
ssh : disable
tftp : enable
ftp : disable

Note: In 4.0.x users are not allowed to disable all the services available in the access policy. For example if five out of the six services are already disabled in an access-policy, the user is not allowed to disable the last enabled service.

In 4.1.x (after saving config in verbose mode and upgrading)

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/6# info

Sub-Context: service

Current Context:

create :
delete : N/A
name : policy6
policy enable : true
mode : allow
precedence : 10
network : 0.0.0.0/0.0.0.0
host : 0.0.0.0



username : none
accesslevel : readOnly
access-strict : false

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/6/service# info

Sub-Context:

Current Context:

http : disable
rlogin : disable
telnet : disable
ssh : disable
snmpv3 : disable
tftp : enable
ftp : disable

Ethernet Routing Switch -8603:3/config/sys/access-policy/policy/6# snmp-group-info

snmpv3-groups :

Group Name Snmp-Model

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks Limited.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>



Ethernet Routing Switch 8600

Software Release 4.1.1.1

1. Release Summary

Release Date: 19-Dec-2006

Purpose: Software patch release to address customer found software issues.

2. Important Notes before Upgrading to This Release

With this release, High Availability CPU mode (HA or hot standby) is now supported within the Ethernet Routing Switch 8600. Nortel recommends using HA-CPU mode as required and based upon known restrictions (reference 4.1.1.0 RN) with release 4.1.1.1 and beyond.

3. Platforms Supported

Ethernet Routing Switch 8600 modules in 8006, 8010, and 8010 co chassis.

Please refer to the following documents for details on the Platforms Supported:

Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E)

Installing Ethernet Routing Switch 8600 Modules. (Part No. 312749-K)

Managing Platform Operations Ethernet Routing Switch 8600 Software Release 4.1 (Part No. 315545-E)

Installing and Maintaining the Ethernet Routing Switch 8000 Series Chassis (Part No.316314-F)

4. Notes for Upgrade

Please see Release Notes for the Ethernet Routing Switch 8600 Series Switch Software Release 4.1.1.0 (Part No. 317177-E) available at:

<http://www.nortel.com/support> (select Ethernet Routing Switch product category for details on how to upgrade your Ethernet Routing Switch 8600)

File Names for This Release

File Name	Module or File Type	File Size (bytes)
p80b4111.img	Boot monitor image	1081431
p80a4111.img	Runtime image	8719766
p80j4111.dld	Run-time image for R modules	1268084



p80c4111.img	3DES	55928
p80c4111.aes	AES	26112
p80a4111.mib	MIB	3319366
p80a4111.mib.zip	MIB (zip file)	531719
p80a4111.md5	md5 checksum file	1488
p80p4111.dld	8600 POS module image	701771
p80t4111.dld	8600 ATM module image	906024
p80m4111.img	SuperMezz image	8817999
<i>Firmware images</i>		
foq267.xsvf	FOQ	5320469
bmc776.xsvf	BMC	2640266
dpc184.xsvf	DPC	2583454
PI_769.xsvf	PIM FPGA	2284578
<i>WSM Images</i>		
wsm1003400_mp.img	WebOS firmware image	845560
wsm1003400_bin.img	WebOS binary	1376256
wsm1003400_boot.img	WebOS boot image	43004

5. Version of Previous Release

Software Version **4.1.1.0**

6. Compatibility

This software release is managed with Java Device Manager (JDM) release **6.0.0.0**

This software release supports the Web Switching Module (WSM) running version WebOS **10.0.34.0**

7. Changes in This Release

Problems resolved in This Release

An ERS8600 with HA-CPU (hot standby) mode enabled, and containing an installed 8683XZR module, will no longer have any 8683XZR ports become inactive after an HA failover has completed. (Q01484430)



The System Configuration parameter "#! flags system-monitor false" will now be displayed as TRUE when issuing the command "show config". Previously the flag displayed a value of false, even though it was operationally true. (Q01485344-01)

An ERS8600 containing an installed 8683XZR module will now show proper configuration options for clock source and framing options for its WAN ports under all conditions . (Q01493665, Q01494672)

Users will now be able to login to an ERS8600 using SSH with Public Key Authentication. (Q01483488-01)

An ERS8600 running OSPF and having an OSPF route with 32 bit mask that overlaps another OSPF route, will no longer become unstable. (Q01491868-02)

If repeated Telnet or FTP sessions are opened against an 8600 In Band and/or out of band IP addresses with quick closures, over the course of time this could lead to system instability. This situation is no longer present with this code release. (Q01462275-01)

The ERS8600 will no longer have any traffic loss when MLT ports with an 'OCP Vendor' GBIC are bounced. (Q01495927)

ERS8600 in a DUAL CPU non-HA mode (warm standby) configuration, with an 8683XZR card present will no longer have the 8683XZR module become inactive when the master CPU is reset. (Q01487003)

When any fdb-filter is configured on an ERS8600, the MAC address record will now be shown within the forwarding table with type as 'mgmt'. Note that these filters still only apply to legacy/classic module ports, but can be used in a mixed chassis configuration. (Q01497977)

8630GBR and 8683XZR cards will no longer become inactive when the master CPU is failed on a switch in HA-CPU (hot standby) mode. (Q01493870).

An ERS8600 switch configured with an 8692 with the SuperMezz CPU option, will no longer become unstable when SSH is disabled. (Q01497587-01)

An ERS8600 will no longer become unstable while R-login sessions are active. (Q01465058-03)

For an ERS8600 in HA mode and having installed an 8683XZR, the value of clock source parameter will no longer get defaulted after having an HA failover. However there is a limitation to this problem which is listed in the "Known Limitations" section of this doc. (Q01520484)

For an ERS8600 running OSPF and having OSPF routes based on type 3 summary LSA, when the destination network of these routes is removed, the routes will get purged out of the routing table in much less time. This problem was resolved in release 4.1.1.0. (Q01409753)

Old Features Removed From This Release

None

8. Outstanding Issues

If the allow-more-specific-route parameter is enabled and the configuration is saved, the value of the parameter is lost after switch reboot. (Q01491654)

During an FPGA firmware upgrade, if the switch or an R-module is reset, the R-module may not get initialized. During these upgrades, the user MUST wait for the returned CLI output, "FPGA UPDATE SUCCESS", and then wait an additional 30 seconds or so, before rebooting or resetting the switch to make the FPGA firmware change active. (Q01501599-01)

In any RSMLT design running L3 IP Multicast, the traffic can stop flowing across the setup upon the re-boot and recovery of a failed switch. (Q01505335)



In a configuration having R-modules and running L3 IP multicast traffic, if the parameter access-diffserv is set to true, the IP Multicast traffic may not be properly sent out the downstream port to the IP Multicast receivers. Do not set this parameter value on R-module ports in a system that requires a L3 IPMC configuration. (Q01479569-01)

RMON traps "Last System Change" & "Ethernet Octets stats" are not recognized by DM and Enterprise Switch Manager (ESM). (Q01142068)

An ERS8600 using OSPF may have the following problems with opaque LSA:

1. When an opaque aware and a non-opaque aware router are mixed, opaque aware should become DR (Designated Router) but in such a case, the ERS8600 will become DR instead.
2. A router should not respond to a LSA type 9 packet if it's 'O' bit is not set but the ERS8600 does.
3. The ERS8600 includes the LSA into its DD packet, which makes it an invalid packet. (Q01509489)

On a switch running cp-limit feature, if redirect-next-hop is configured using an ACL filter and global-action is set to count or mirror-count the port having ingress traffic is shut down if the specified next-hop is unreachable. Do not configure ACL redirect filters, with additional optional action of count or mirror-count (Q01493502, Q01492644)

R-module ports currently do not properly handle tagged multicast traffic which must egress on the same physical port that the traffic ingresses on. This affects any and all multicast based traffic or applications, of which Microsoft NLB is one example. To workaround this issue, make sure that the client and server are connected to the ERS8600 by different physical ports – different L2 switches in a SMLT environment, for example. Additionally, this configuration on legacy module ports will provide a working environment. (Q01511596)

On an ERS8600 installed with 8630 GBR card, if an egress queue template is dynamically applied without a switch reboot, the 8630GBR ports can get into a condition whereby individual port egress queues seize up and the port is no longer able to transmit any frames. If changes to an egress queue are required, the user MUST reboot or reset the switch after the changes are made. In a future code release, a warning message will be generated to the user for this activity. (Q01450460-01)

For an ERS8600 connected to an ES470 with SFFD being used on the connection, and if the connection is multiple links to an 8630GBR module, failure or pulling out of one of the links, may cause the remaining links to flap. Instead configure the connection between the ERS8600 and ES470 to use VLACP instead of SFFD, which may require a code change on the ES470. (Q01464922).

For an ERS8600 configured with VRRP, the show config or save config commands will fail and CPU utilization will increase to 100% on removing the VRRP IP address if the VRRP IP is also being used as a DHCP agent. VRRP addresses should not be used as DHCP relay agents. Instead the physical IP interface address associated with the VLAN should be used. (Q01473196-01)

Trying to accessing a switch with an installed WSM, running with version 10.0.34.0, from JDM with version 6.0.3, the following error message is displayed:
ERROR: Cards in slots: slot# is running WebOS versions that incompatible with this version of Device Manager. Please upgrade these cards to a later patch version." The message is erroneous, and can be ignored. (Q01491448)

In any SMLT design for the ERS8600, ARP entries on clients connected to the SMLT client L2 switch may be learned on the IST instead of the SMLT links after the MAC ages out. This can lead to intermittent user connectivity issues. If seen, a flush of the ARP table (config ip arp flush), will correct the situation, at least temporarily. Additionally setting the per VLAN FDB age out timer for all affect VLANs, to 21601 (1 second higher than default system ARP age out timer), may assist as well. (Q01476117)

On an ERS8600 installed with ROHS compliant 8005AC power supplies DS1405012-E5, P/N 321411, the total available power is not calculated correctly. Due to this the power supplies are not utilized fully. (Q01495801)

Upon issuing the command "show port info state" on an ERS8600, current time/date is displayed for a port in a down state. The correct behavior should be to display the date/time of the last state change. (Q01474469-01)



On an ERS8600, when an MLT is created using a null string for the name of MLT, any further attempt to create a new MLT with CLI will fail with the following error message: "Error: MLT invalid mlt name". Do not create MLTs with assigning them some name. (Q01478418-02)

An ERS8600 installed with R-modules, enabling the Tagged-Frames-Discard on any port of the R-module will stop the IPV4/IPV6 bridging function of any protocol based VLANs associated with the port. Do not use this configuration option with ports associated with any protocol based VLAN. (Q01427096)

On an ERS8600 with Dual SF/CPU modules installed, the switch fabric utilization LED behavior may not always be shown correctly. (Q01497340)

On an ERS8600 with multiple WSM modules installed, and with a large configuration on the WSM modules, when the switch is booted, some of the WSM cards might not be manageable except by the direct console connection, but the module will pass traffic as normal. (Q01473364-01)

When an ERS8600 running VRRP is connected to a client through a layer 2 forwarding switch, under certain conditions the 8600 may not reply for the ARPs sent for the VRRP IP address by the client. A ping of the end-station client from the 8600, will resolve the situation. (Q01502287)

When setting dst-end to desired date and time, DST (Daylight Savings Time) does not take effect, as changes do not take effect once the time is reached. (Q01501308)

When the switch CPU utilization reaches near constant 100% on an ERS8600, the SF/CPU module may hang and will not fail-over for 1 to 2 minutes. Any sustained and constant high CPU utilization situations should always be investigated, prior to the event causing any network disruption. (Q01374277)

Under certain situations, for an ERS8600 with HA-CPU mode (hot standby) enabled, and with having multiple STGs, entries in the STG topology table may be lost after an HA-CPU failover. (Q01479548-02)

An ERS8600 with R-module installed, JDM does not display Egress Queue Set Queue 0. (Q01490598-01)

On a switch running PIM-SM, changing of activity-chk-interval parameter should not be allowed. Currently, the ERS8600 allows it. Users should use just the default value of 210 for this parameter. (Q01464950-01; Q01441360-01)

9. Known Limitations

Please refer to the "Known Limitations" section of the Ethernet Routing Switch 8600 Software Release 4.1.1.0 (Part Number 317177-E).

In certain cases, but only with 8692 and SuperMezz SF/CPU and with the Mezz option enabled, upon a power cycle the chassis may come with slot 6 as the master, versus slot 5, even when master configuration option is set to slot 5. To make slot 5 the master, a reset or re-boot on slot 6 must take place. (Q01522492)

When using the SuperMezz option, the mezz image (p80mxxxx.img) must match the previously load boot image (p80bxxxx.img) or else the SuperMezz will not load/boot. (Q01522508)

10. Documentation Corrections

1. When replacing or re-inserting any module with an ERS8600 chassis, the user MUST wait at least 30 seconds, before putting in the replacement module, or for reseating the existing module. This requirement applies equally to the SF/CPU, as it does to I/O modules. (Q01489847)



2. The following documentation change is associated with Page 70 of the Configuring IP Multicast Routing Protocols manual (314719-D):

The most recommended L3 IPMC implementation with PIM-SM on and ERS8600 is to configure a candidate RP with a multicast group address and mask that is as close to the multicast source as possible. The candidate BSR should be configured on a router that is central to the entire candidate RP's. (Q01510496)

As well, a candidate RP can provide service to all potential multicast groups (224.0.0.0 – 239.255.255.255) if its GroupAddress value is set to 224.0.0.0 with an associated GroupMask of 240.0.0.0 (see page 240 of the above noted manual).

3. ERS8600 manual Using Diagnostic Tools (317359-D) on page 42 states:

Configuration example: Creating mirror-by-mirror entries, when it should read:

Configuration example: Creating mirror-by-port entries.

4. In the release notes for Ethernet Routing Switch 8600 Software Release 4.1.1. 0, the file sizes of p80c4110.img (3DES image) and p80c4110.aes (AES image) are given wrongly (sizes are swapped). The correct sizes in bytes are:

p80c4110.img: 55928

p80c4110.aes: 26112

5. When setting multiple ports of the 8683XZR module to WAN-mode, the frame Type and CkMode have to be set the same for all the ports configured for WAN-mode. Mixed settings for multiple WAN-mode enabled ports is not allowed.

6. When upgrading FPGA firmware images, the user should (current Upgrade Guide does not call this out exactly; next revision will):

a. Never download an image or multiple images to different modules at the same time. Only download one image to one module at a time. FPGA firmware images can be downloaded with no impact on system operation. The only action that affects system operation is the switch reboot which is required to make all FPGA firmware updates take affect.

b. The user should always wait 30 seconds after the return to the telnet or console screen of the message "FPGA UPDATE SUCCESS", before issuing a system reboot for the FGPA upgrade(s) to take affect.

A more detailed log message will be added into a future code release to warn users of this required action. (Q01509339)

Copyright © 2006 Nortel Networks - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globe mark, and Ethernet Routing Switch 8100/8300/8600 are trademarks of Nortel Networks.

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 Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at: <http://www.nortel.com/support>