



Nortel Ethernet Routing Switch 8600

# Release Notes — Software Release 5.1

Release: 5.1  
Document Revision: 03.02

[www.nortel.com](http://www.nortel.com)

---

NN46205-402

Nortel Ethernet Routing Switch 8600  
Release: 5.1  
Publication: NN46205-402  
Document release date: 4 June 2009

Copyright © 2008-2009 Nortel Networks. All Rights Reserved.

While the information in this document is believed to be accurate and reliable, except as otherwise expressly agreed to in writing NORTEL PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED. The information and/or products described in this document are subject to change without notice.

Nortel, Nortel Networks, the Nortel logo, and the Globemark are trademarks of Nortel Networks.

All other trademarks are the property of their respective owners.

---

# Contents

---

|                                                                      |           |
|----------------------------------------------------------------------|-----------|
| <b>Software license</b>                                              | <b>7</b>  |
| <b>New in this release</b>                                           | <b>11</b> |
| Features 11                                                          |           |
| New hardware and software 11                                         |           |
| Other changes 11                                                     |           |
| Considerations for upgrades from pre-5.0.1.0 release code 11         |           |
| Changes to upgrade procedures for Release 5.1 software licensing 12  |           |
| 60 day temporary license 12                                          |           |
| File names for upgrade 12                                            |           |
| Supported scaling capabilities 12                                    |           |
| RSMLT considerations 12                                              |           |
| Known issues from release 5.1 12                                     |           |
| Customer service chapter 12                                          |           |
| <b>Introduction</b>                                                  | <b>13</b> |
| <b>Important notices and new features</b>                            | <b>15</b> |
| New features in Release 5.1 16                                       |           |
| Multicast virtualization for VRF-Lite (IGMP, PIM-SM, and PIM-SSM) 18 |           |
| Multicast Source Discovery Protocol (MSDP) 18                        |           |
| IGMPv3 backward compatibility 18                                     |           |
| Multicast static IP route (mroute) 18                                |           |
| Static RP enhancement for PIM-SM 18                                  |           |
| Standards-based MIBs for multicast 19                                |           |
| Multicast routing process statistics 19                              |           |
| IPv6 support for FTP, TFTP, and rlogin 19                            |           |
| TACACS+ 19                                                           |           |
| Bidirectional Forwarding Detection (BFD) 19                          |           |
| LACP minimum link 19                                                 |           |
| DDI software support 20                                              |           |
| Time zone offset enhancement 20                                      |           |
| ping and traceroute source IP enhancement 20                         |           |
| SMLT/RSMLT operational improvements 20                               |           |
| IPFIX enhancement 25                                                 |           |

---

|                                                                              |    |
|------------------------------------------------------------------------------|----|
| Nortel Automatic QoS                                                         | 25 |
| New operation for corrupt configuration loading                              | 28 |
| Core dumps                                                                   | 28 |
| Nortel qualified DWDM XFPs                                                   | 30 |
| Support for 100Base-FX SFP (AA1419074)                                       | 30 |
| Support for new re-spin version of 8692SF/CPU (no part number or PEC change) | 31 |
| Software licensing                                                           | 31 |
| File names for this release                                                  | 32 |
| Important information and restrictions                                       | 35 |
| SuperMezz, CPU memory, and upgrades                                          | 36 |
| Nonsupported modules for Release 5.0 and later                               | 39 |
| I/O module considerations                                                    | 39 |
| Mixed-module MLT/LAG support for Release 5.0 and later                       | 39 |
| Fixes from previous releases                                                 | 40 |
| Supported upgrade paths                                                      | 41 |
| General upgrade considerations                                               | 41 |
| Upgrade considerations for Release 5.x                                       | 41 |
| Upgrade considerations: DOSFS with upgrades from pre-Release 5.0             | 41 |
| Upgrade considerations: Power Management                                     | 42 |
| Disabling power and cooling management                                       | 43 |
| Upgrade considerations: IST                                                  | 45 |
| Changes to options of license file installation procedures for Release 5.1   | 46 |
| Considerations for upgrades from pre-5.0.1.0 release code                    | 47 |
| Supported software and hardware scaling capabilities                         | 52 |
| Hardware and software compatibility                                          | 57 |
| Module and mode compatibility                                                | 63 |
| Module type and feature compatibility                                        | 63 |
| Module and chassis compatibility and performance considerations              | 64 |
| High Performance chassis                                                     | 65 |
| Switch clustering topologies and interoperability with other products        | 66 |
| SF/CPU protection and loop prevention compatibility                          | 66 |
| Switch behavior during boot cycle and redundant configuration files          | 67 |
| Configuring primary, secondary, and tertiary boot sources                    | 68 |
| MPLS considerations                                                          | 70 |
| SNMP considerations                                                          | 70 |
| VLACP considerations                                                         | 71 |
| DVMRP considerations                                                         | 71 |
| SMLT considerations                                                          | 71 |
| RSMLT considerations                                                         | 72 |
| MSDP considerations                                                          | 73 |
| High Availability mode considerations                                        | 73 |
| HA mode and short timers                                                     | 74 |
| PCAP considerations                                                          | 74 |

---

|                                                |    |
|------------------------------------------------|----|
| WSM considerations                             | 75 |
| Enterprise Switch Manager (ESM) considerations | 75 |
| 60 day trial license                           | 75 |
| Advanced filter guidelines                     | 75 |
| Supported standards, RFCs, and MIBs            | 76 |
| Supported traps and notifications              | 76 |
| Device Manager installation requirements       | 76 |
| Windows                                        | 77 |
| Solaris                                        | 77 |
| Linux                                          | 77 |
| Software Release 3.7 support                   | 77 |

---

## **Resolved issues**

**79**

|                                                |    |
|------------------------------------------------|----|
| Layer 2 switching resolved issues              | 79 |
| Switch management resolved issues              | 80 |
| Platform resolved issues                       | 81 |
| IP Unicast resolved issues                     | 83 |
| MLT/SMLT resolved issues                       | 84 |
| CLI and NNCLI resolved issues                  | 85 |
| BGP resolved issues                            | 86 |
| Multicast resolved issues                      | 86 |
| Quality of Service and filters resolved issues | 87 |
| Device Manager resolved issues                 | 87 |

---

## **Known issues and limitations**

**89**

|                                             |     |
|---------------------------------------------|-----|
| Release 5.1 known issues                    | 89  |
| Platform known issues                       | 89  |
| Switch management known issues              | 92  |
| Layer 2 known issues                        | 93  |
| MLT/SMLT known issues                       | 94  |
| Unicast routing known issues                | 94  |
| Multicast routing known issues              | 96  |
| CLI and NNCLI known issues                  | 97  |
| Quality of Service and filters known issues | 97  |
| Device Manager known issues                 | 98  |
| Release 5.1 known limitations               | 99  |
| Platform limitations                        | 99  |
| Switch management limitations               | 102 |
| Layer 2 limitations                         | 103 |
| MLT/SMLT limitations                        | 104 |
| Unicast routing limitations                 | 105 |
| Multicast routing limitations               | 105 |
| QoS and filters limitations                 | 105 |
| Device Manager limitations                  | 106 |
| MIB limitations                             | 106 |

---

---

|                                   |            |
|-----------------------------------|------------|
| <b>Customer service</b>           | <b>107</b> |
| Updated versions of documentation | 107        |
| Getting help                      | 107        |
| Express Routing Codes             | 107        |
| Additional information            | 108        |
| <b>Index</b>                      | <b>109</b> |

---

---

## Software license

---

This section contains the Nortel Networks software license.

### **Nortel Networks Inc. software license agreement**

This Software License Agreement ("License Agreement") is between you, the end-user ("Customer") and Nortel Networks Corporation and its subsidiaries and affiliates ("Nortel Networks"). PLEASE READ THE FOLLOWING CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

"Software" is owned or licensed by Nortel Networks, its parent or one of its subsidiaries or affiliates, and is copyrighted and licensed, not sold. Software consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings or pictures) and related licensed materials including all whole or partial copies. Nortel Networks grants you a license to use the Software only in the country where you acquired the Software. You obtain no rights other than those granted to you under this License Agreement. You are responsible for the selection of the Software and for the installation of, use of, and results obtained from the Software.

**1. Licensed Use of Software.** Nortel Networks grants Customer a nonexclusive license to use a copy of the Software on only one machine at any one time or to the extent of the activation or authorized usage level, whichever is applicable. To the extent Software is furnished for use with designated hardware or Customer furnished equipment ("CFE"), Customer is granted a nonexclusive license to use Software only on such hardware or CFE, as applicable. Software contains trade secrets and Customer agrees to treat Software as confidential information using the same care and discretion Customer uses with its own similar information that it does not wish to disclose, publish or disseminate. Customer will ensure that anyone who uses the Software does so only in compliance with the terms

of this Agreement. Customer shall not a) use, copy, modify, transfer or distribute the Software except as expressly authorized; b) reverse assemble, reverse compile, reverse engineer or otherwise translate the Software; c) create derivative works or modifications unless expressly authorized; or d) sublicense, rent or lease the Software. Licensors of intellectual property to Nortel Networks are beneficiaries of this provision. Upon termination or breach of the license by Customer or in the event designated hardware or CFE is no longer in use, Customer will promptly return the Software to Nortel Networks or certify its destruction. Nortel Networks may audit by remote polling or other reasonable means to determine Customer's Software activation or usage levels. If suppliers of third party software included in Software require Nortel Networks to include additional or different terms, Customer agrees to abide by such terms provided by Nortel Networks with respect to such third party software.

**2. Warranty.** Except as may be otherwise expressly agreed to in writing between Nortel Networks and Customer, Software is provided "AS IS" without any warranties (conditions) of any kind. NORTEL NETWORKS DISCLAIMS ALL WARRANTIES (CONDITIONS) FOR THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OF NON-INFRINGEMENT. Nortel Networks is not obligated to provide support of any kind for the Software. Some jurisdictions do not allow exclusion of implied warranties, and, in such event, the above exclusions may not apply.

**3. Limitation of Remedies.** IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER'S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN ADVISED OF THEIR POSSIBILITY. The forgoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

#### **4. General**

1. If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer

software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

2. Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.
3. Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.
4. Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.
5. The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.
6. This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.



---

## New in this release

---

The following sections detail what's new in *Nortel Ethernet Routing Switch 8600 Release Notes — Software Release 5.1* (NN46205-402) for Release 5.1.

- [“Features” \(page 11\)](#)
- [“Other changes” \(page 11\)](#)

### Features

See the following sections for information about feature changes.

#### **New hardware and software**

Release 5.1 adds many new features and devices. For more information about these, see [“New features in Release 5.1” \(page 16\)](#).

### Other changes

See the following sections for information about changes that are not feature-related.

#### **Considerations for upgrades from pre-5.0.1.0 release code**

When upgrading a pair of IST Core peer switches, special care must be taken due to SMLT architectural change. Unlike previous upgrades, when moving from pre-5.0.1.0 code to any 5.0.1.0 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.

For more information, see [“Considerations for upgrades from pre-5.0.1.0 release code” \(page 47\)](#).

### **Changes to upgrade procedures for Release 5.1 software licensing**

After you upgrade to release 5.1, you must specify the name and location of your license file in the boot configuration file. If you do not specify the name and location of your license file, you can encounter issues with your licensed features. See [“Changes to options of license file installation procedures for Release 5.1” \(page 46\)](#).

The software licensing section lists all licensed features contained in software Release 5.1. See [“Software licensing” \(page 31\)](#).

### **60 day temporary license**

You are provided a 60 day trial period of Ethernet Routing Switch 8600, during which you have access to all features. The message Trial Period for Automatic Premier Feature usage will expire in ## days appears every 24 hours during the trial period. If any valid license is loaded on the switch at any time, no messages are recorded. For more information, see [“60 day trial license” \(page 75\)](#).

### **File names for upgrade**

File names are updated; see [“File names for this release” \(page 32\)](#).

### **Supported scaling capabilities**

The supported scaling capabilities table is updated to include DHCP relay scaling information (CR Q01949340), as well as information about OSPF instances, RIP instances, and PIM instances for each switch. See [“Supported software and hardware scaling capabilities” \(page 52\)](#).

### **RSMLT considerations**

A section on RSMLT considerations is added. See [“RSMLT considerations” \(page 72\)](#).

### **Known issues from release 5.1**

A known issues list is updated for release 5.1. See [“Known issues and limitations” \(page 89\)](#).

### **Customer service chapter**

A customer service chapter is now added to this document. See [“Customer service” \(page 107\)](#).

# Introduction

---

This document describes new features and known and fixed issues for Ethernet Routing Switch 8600 Software Release 5.1. Use this document to help you optimize the functionality of your Ethernet Routing Switch 8600.

For information about how to upgrade your version of Device Manager, see *Nortel Ethernet Routing Switch 8600 User Interface Fundamentals* (NN46205-308).

## Navigation

- [“Important notices and new features” \(page 15\)](#)
- [“Resolved issues” \(page 79\)](#)
- [“Known issues and limitations” \(page 89\)](#)
- [“Customer service” \(page 107\)](#)



---

## Important notices and new features

---

This section describes the supported and unsupported hardware and software features in the Ethernet Routing Switch 8600 Software Release 5.1, describes new features, and provides important information for this release.

### Navigation

- [“New features in Release 5.1” \(page 16\)](#)
- [“Software licensing” \(page 31\)](#)
- [“File names for this release” \(page 32\)](#)
- [“Important information and restrictions” \(page 35\)](#)
- [“Supported software and hardware scaling capabilities” \(page 52\)](#)
- [“Hardware and software compatibility” \(page 57\)](#)
- [“Module and mode compatibility” \(page 63\)](#)
- [“Module type and feature compatibility” \(page 63\)](#)
- [“Module and chassis compatibility and performance considerations” \(page 64\)](#)
- [“Switch clustering topologies and interoperability with other products” \(page 66\)](#)
- [“SF/CPU protection and loop prevention compatibility” \(page 66\)](#)
- [“Switch behavior during boot cycle and redundant configuration files” \(page 67\)](#)
- [“MPLS considerations” \(page 70\)](#)
- [“SNMP considerations” \(page 70\)](#)
- [“VLACP considerations” \(page 71\)](#)
- [“DVMRP considerations” \(page 71\)](#)
- [“SMLT considerations” \(page 71\)](#)
- [“MSDP considerations” \(page 73\)](#)

- “High Availability mode considerations” (page 73)
- “PCAP considerations” (page 74)
- “Enterprise Switch Manager (ESM) considerations” (page 75)
- “60 day trial license” (page 75)
- “WSM considerations” (page 75)
- “Advanced filter guidelines” (page 75)
- “Supported standards, RFCs, and MIBs” (page 76)
- “Supported traps and notifications” (page 76)
- “Device Manager installation requirements” (page 76)
- “Software Release 3.7 support” (page 77)

## **New features in Release 5.1**

For 5.1, the Ethernet Routing Switch 8600 software supports:

- Multicast virtualization for VRF-Lite (IGMP, PIM-SM, and PIM-SSM)
- Multicast Source Discovery Protocol (MSDP)
- IGMPv3 backward compatibility
- Multicast Static IP Route (mroute)
- Static RP enhancement for PIM-SM
- Standards-based MIBs for multicast
- Multicast routing process statistics
- IPv6 support for FTP, TFTP, and rlogin
- TACACS+
- Bidirectional Forwarding Detection (BFD)
- Link Aggregation Control Protocol (LACP) Minimum Link
- DDI software support
- Time zone offset enhancement
- ping and traceroute source IP enhancement
- SMLT/RSMLT operational improvements
- IPFIX enhancement
- Nortel Automatic QoS
- New operation for corrupt configuration loading
- Core dumps

New hardware includes:

- Nortel qualified DWDM XFPs
- Support for 100Base-FX SFP (AA1419074)
- Support for new re-spin version of 8692SF/CPU (no part number or PEC change)

### **New features in Release 5.1 navigation**

- [“Multicast virtualization for VRF-Lite \(IGMP, PIM-SM, and PIM-SSM\)” \(page 18\)](#)
- [“Multicast Source Discovery Protocol \(MSDP\)” \(page 18\)](#)
- [“IGMPv3 backward compatibility ” \(page 18\)](#)
- [“Multicast static IP route \(mroute\)” \(page 18\)](#)
- [“Static RP enhancement for PIM-SM” \(page 18\)](#)
- [“Standards-based MIBs for multicast” \(page 19\)](#)
- [“Multicast routing process statistics” \(page 19\)](#)
- [“IPv6 support for FTP, TFTP, and rlogin” \(page 19\)](#)
- [“TACACS+” \(page 19\)](#)
- [“Bidirectional Forwarding Detection \(BFD\)” \(page 19\)](#)
- [“LACP minimum link” \(page 19\)](#)
- [“DDI software support” \(page 20\)](#)
- [“Time zone offset enhancement” \(page 20\)](#)
- [“ping and traceroute source IP enhancement” \(page 20\)](#)
- [“SMLT/RSMLT operational improvements” \(page 20\)](#)
- [“IPFIX enhancement” \(page 25\)](#)
- [“Nortel Automatic QoS” \(page 25\)](#)
- [“New operation for corrupt configuration loading” \(page 28\)](#)
- [“Core dumps” \(page 28\)](#)
- [“Nortel qualified DWDM XFPs” \(page 30\)](#)
- [“Support for 100Base-FX SFP \(AA1419074\)” \(page 30\)](#)
- [“Support for new re-spin version of 8692SF/CPU \(no part number or PEC change\)” \(page 31\)](#)

### **Multicast virtualization for VRF-Lite (IGMP, PIM-SM, and PIM-SSM)**

Release 5.1 adds multicast routing support for the Virtual Router Forwarding (VRF) Lite feature. You can use VRF Lite to emulate many virtual routers with one router. In this release, VRF Lite virtualizes the following multicast routing protocols: IGMP, Protocol Independent Multicast - Source Mode (PIM-SM), and Protocol Independent Multicast - Source-Specific Multicast (PIM-SSM). For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Multicast Routing Protocols* (NN46205-501).

### **Multicast Source Discovery Protocol (MSDP)**

MSDP connects multiple IP version 4 Protocol Independent Multicast Sparse-Mode (PIM-SM) domains. Rendezvous points (RPs) in different domains use MSDP to discover and distribute multicast sources for a group. Each PIM-SM domain uses an independent RP and does not depend on RPs in other domains. For more information about MSDP, see *Nortel Ethernet Routing Switch 8600 Configuration — Multicast Source Discovery Protocol* (NN46205-524).

### **IGMPv3 backward compatibility**

Beginning with Release 5.1, Internet Group Management Protocol version 3 (IGMPv3) for Protocol Independent Multicast-Source Specific Multicast (PIM-SSM) is backward compatible with IGMPv1 and IGMPv2. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Multicast Routing Protocols* (NN46205-501).

### **Multicast static IP route (mroute)**

Release 5.1 introduces a new static IP route table to separate the paths for unicast and multicast streams. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Multicast Routing Protocols* (NN46205-501).

### **Static RP enhancement for PIM-SM**

With static RP enabled, the Ethernet Routing Switch detects RP failure based on IGP convergence and, more specifically, on the removal of the route to the RP from the routing table. With the route to the failed RP removed, the Ethernet Routing Switch can fail over to an alternate static RP.

However, when a default route is injected into the routing table, that default route still appears as an active route to the failed RP. Therefore, in this case, the switch does not failover to the alternate RP.

To work around this issue, in Release 5.1, you can configure the lookup for static RP to be chosen from the specific route rather than the best route. In this case, when the route to the active RP fails, the switch no longer interprets the default route as a valid route for RP purposes, and so fails over to the alternate RP.

### **Standards-based MIBs for multicast**

Release 5.1 replaces proprietary multicast management information base (MIB) definitions with standardized MIB definitions. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Multicast Routing Protocols* (NN46205-501).

### **Multicast routing process statistics**

To provide additional troubleshooting information, you can enable the collection and display of multicast routing process statistics. These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on VRF and multicast group classification. For more information, see: *Nortel Ethernet Routing Switch 8600 Troubleshooting* (NN46205-703).

### **IPv6 support for FTP, TFTP, and rlogin**

With release 5.1, the Ethernet Routing Switch 8600 supports IPv6 addresses for FTP, TFTP, and rlogin access.

### **TACACS+**

Ethernet Routing Switch 8600 Series release 5.1 supports the Terminal Access Controller Access Control System plus (TACACS+) client. TACACS+ is a security application implemented as a client/server-based protocol that provides centralized validation of users attempting to gain access to a router or network access server. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Security* (NN46205-601).

### **Bidirectional Forwarding Detection (BFD)**

Release 5.1 provides support for Bidirectional Forwarding Detection (BFD). Bidirectional Forwarding Detection (BFD) is a simple Hello protocol used between two peers. In BFD, each peer system periodically transmits BFD packets to each other. If one of the systems does not receive a BFD packet after a certain period of time, the system assumes that the link or other system is down. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration -- BGP Services* (NN46205-510) and *Nortel Ethernet Routing Switch 8600 Configuration -- IP Routing* (NN46205-523).

### **LACP minimum link**

Release 5.1 provides support for LACP minimum link. The Minimum Link parameter defines the minimum number of active links required for a LAG to remain in the forwarding state. Use the Minimum-Link (MinLink) feature

so that when the number of active links in a LAG is less than the MinLink parameter, the entire LAG is declared down. Prior to MinLink support, a LAG was always declared up if one physical link of the LAG was up. For more information, see *Nortel Ethernet Routing Switch 8600 Configuration — Link Aggregation, MLT, and SMLT* (NN46205-518).

### **DDI software support**

Release 5.1 adds support for Digital Diagnostic monitoring. Digital Diagnostic monitoring provides a user with critical information concerning the status of the transmitted and received signals. This approach allows for better fault isolation and error detection. For more information, see *Nortel Ethernet Routing Switch 8600 Performance Management* (NN46205-704).

### **Time zone offset enhancement**

The format for the time zone offset commands are changed to support entry of the time offset as a function of hours and minutes rather than just minutes. For more information, see *Nortel Ethernet Routing Switch 8600 Commissioning* (NN46205-319)

### **ping and traceroute source IP enhancement**

Release 5.1 allows you to specify the source IP address for ping and traceroute using `ping <target-ip> source <source-ip>` and `traceroute <target-ip> source <source-ip>` commands.

### **SMLT/RSMLT operational improvements**

For previous SMLT operation, bringing the SMLTs Up/Down triggered the flushing of the entire MAC/ARP belonging to the SMLTs in both the IST core peer switches. Flushing of the MAC addresses then causes the dependent ARP 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 are flushed, during the re-learning period the exception (learning) packets will also be continuously forwarded to the CP, thereby increasing the CPU load. This will further slow-down the SMLT re-convergence as well as the h/w record re-programming.

Since the traffic flow 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 also running with many multicast streams, as multicast streams often need to be forwarded to CPU for learning, thereby also increasing CPU load. The SMLT changes 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. The SMLT/RSMLT associated links will not become active upon a boot-up until the IST is completely up, and a 'sync' has occurred between the two Core IST Peer Switches.

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.

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. These messages formats will be seen as below, under the various situations noted.

**CAUTION**

Read "[Considerations for upgrades from pre-5.0.1.0 release code](#)" ([page 47](#)) before performing any upgrade. For any pair of IST Core Switch Cluster upgrades, to upgrade to 5.0.1.0 or later code from any 5.0.0.x release, the user must first upgrade to the 5.0.0.2 (released code can be found on web at [www.nortel.com/support](http://www.nortel.com/support)). This ONLY applies to upgrading a pair of IST Core switches (does not apply to single switches or non-SMLT Core pairs of ERS 8600 switches) and only if the user is currently running any 5.0.0.x code release. This does not apply to upgrades from any 4.1.x release. (CR Q1916209). To upgrade from 5.0.0.x to 5.0.0.2, see *Nortel Ethernet Routing Switch 8600 Upgrades — Software Release 5.0* (NN46205-400). The Upgrade Consideration section of this document applies to upgrades to 5.0.1.0 only.

**CAUTION**

On any switch, whether or not it is part of a Core IST Cluster or a single switch, if a downgrade is required from 5.0.1.0 or later to some 5.0.0.x code, **ONLY** downgrade to 5.0.0.2 code. The prior saved 5.0.0.x config file will work with 5.0.0.2 code. If this is not done, there is a small potential to corrupt the boot image on R/RS modules; so if the system has no R/RS modules, then this is not a consideration. Another option, instead of going to 5.0.0.2 code for the downgrade, is to remove all R/RS modules from the chassis prior to the downgrade, and then re-insert one slot at a time, after the downgrade of the SF/CPU module (most likely some form of an 8692SF/CPU) has been completed. (CR Q01971801)

**Table 1**  
**SMLT informational message examples**

| Case                                                                                                                                                                                                                                                                                                                                                                                                           | Message                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch running SMLT is reset. Upon switch coming up online only one message is displayed irrespective of the number of SMLTs.                                                                                                                                                                                                                                                                                  | <pre>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</pre>                                                                                                                                                                                                                            |
| System is up and running but SMLT UP event 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. To determine which specific SMLT IDs are affect, look for the SMLT ID down/up log messages.                                                                                                                   | <pre>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 SMLTMAC /ARP Sync is Complete :  Peer can now be used normally</pre> |
| 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, or maybe a mis-configuration. If the system that is reset is requesting sync, it will keep all the ports locked down until the IST comes up properly and sync has occurred. We display the error message after 5 minutes. | <pre>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 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</pre>                                                                                                                     |

**Additional SMLT considerations**

Software Release 5.1 does not support PIM Multicast Border Router (MBR) functionality over SMLT.

Nortel does not support an additional redundant IST\_MLT group between two IST peers. Use of a second MLT connection (non-IST) is supported, but such designs are not recommended.

For an IPVPN-lite configuration, where an edge 8600 Cluster is configured to use a SMLT configuration toward the core (most likely square or full-mesh RSMLT), SMLT fast failover cannot always be guaranteed for this portion of the network. (CR Q2016922)

In a large-scale SMLT network (that is, one with more than 10 000 MAC addresses), Nortel strongly recommends that you operate the system with M, or preferably, R and RS modules. Further, if you use R or RS modules only, then you should enable R mode. This ensures that the system has enough memory and record space available for proper switch operation.

**Note 1:** All tested failover times for SMLT/RSMLT operation are sub-second. With this release, all tested fail-back or recovery times have been improved, especially for very large scaled environments, to be within 3 seconds. This worst-case value is used in order to provide required recovery time for converged networks, especially IP phone operation. These values are for unicast traffic only; this includes voice traffic. Not all IP Multicast failover or fail-back/recovery situations can provide such times today, as many situations depend on the IPMC protocol itself recovering. For best IPMC recovery in SMLT/RSMLT designs, the use of static RPs for PIM-SM is recommended, with the same CLIP address assigned to both Core IST Peers within the Cluster, and to all switches within a full-mesh or square configuration. Failover or fail-back/recovery times for any situations that involve high-layer protocols, such as OSPF, BGP, etc., can not always be guaranteed. Reference the Network Design Guide for your specific code release for recommendations on best practices, in order to achieve best results. The 3 second recovery time is supported in Nortel SMLT/RSMLT designs which utilize VLACP. Designs that do not use VLACP or consist of a mixed vendor environment cannot be guaranteed to have a 3 second fail-back or recovery time. Nortel continues to support SMLT designs with non-Nortel devices, but recovery times for these designs cannot always be expected to be within 3 seconds.

**Note 2:** 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 increase 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 ( $5 \times 500 = 2.5$  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 the associated remote connection down, but since the 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 cause a drop of 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. The same applies to LACP configurations for 1 port MLTs.

**Note 3:** If using VRRP with SMLT, users **MUST** 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. An alternative method, which allows scaling to maximum number of IP VLANs, is to use RSMLT designs instead.

**Note 4:** For any SMLT design, for either RSMLT L2 Edge VLANs or L2 SMLT VLANs using VRRP, 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.

**Note 5:** There are now new and improved informational and warning log messages associated with SMLT. Some messages will be generated for normal SMLT failovers, but additionally these messages can assist with troubleshooting any abnormal SMLT activity. Under

abnormal conditions, please contact Nortel Technical Support for further assistance, and provide the exact log messages generated.

### IPFIX enhancement

With Release 5.1, IPFIX and advanced filters allow collection of IP flow information for ACLs that have a deny action. For more information, see *Nortel Ethernet Routing Switch 8600 Performance Management* (NN46205-704).

### Nortel Automatic QoS

The Nortel Automatic QoS feature allows Nortel data products to better support Nortel Converged Voice deployments (VoIP) by automatically recognizing the DSCP values that Nortel Voice applications use, and associating these DSCP values with the proper egress queues. Without Nortel Automatic QoS support, you may need to manually configure the DSCP values on the Ethernet Routing Switch and map them to the appropriate queues. With Nortel Automatic QoS enabled, manual DSCP-to-queue mapping is not required.

The following table shows various traffic types mapped to the standard DSCP values, the Nortel Automatic QoS DSCP values, and their associated queues.

**Table 2**  
**Nortel Automatic QoS DSCP Values**

| Traffic Type              | Standard DSCP Value | Old Queue Value | Nortel Automatic QoS DSCP Value (hex/decimal) | New Queue Value |
|---------------------------|---------------------|-----------------|-----------------------------------------------|-----------------|
| VoIP Data (Premium)       | 0x2E (46)<br>EF     | 6               | 0x2F (47)                                     | 6               |
| VoIP Signaling (Platinum) | 0x28 (40)<br>CS5    | 5               | 0x29 (41)                                     | 5               |
| Video (Platinum)          | 0x22 (34)<br>AF41   | 5               | 0x23 (35)                                     | 5               |
| Streaming (Gold)          | 0x1A (26)<br>SF31   | 4               | 0x1B (27)                                     | 4               |

For proper functioning of the feature, you must enable Nortel Automatic QoS on the Ethernet Routing Switch and on the associated Nortel Voice application.

Nortel Auto QoS is supported on the following Nortel voice and data products:

- **Ethernet Routing Switch 4500**

- Release 5.2
- Edge with Nortel Automatic QoS mixed or pure mode
- **Ethernet Routing Switch 5000**
  - Release 6.0
  - Edge with Nortel Automatic QoS mixed or pure mode
- **Ethernet Routing Switch 8300**
  - Release 4.2
  - Nortel Automatic QoS core only
- **Ethernet Routing Switch 8600**
  - Release 5.1
  - Nortel Automatic QoS core only
- **CS1000**
  - Nortel Automatic QoS supported in Element Manager
  - Release 5.5
  - Patch MPLR26485 is required
- **CS2100**
  - SE10
  - Edge with Nortel Automatic QoS supported in Element Manager
- **BCM 50, SRG 50, and BCM450**
  - BCM50/SRG50 requires minimum of Release 3.0 software with Smart Update BCM050.R300.SU.System-115 or later
  - BCM450 requires minimum of Release 1.0 software with Smart Update BCM450.R100.SU.System-003 or later

For configuration information for these products, see *Nortel Automatic QoS Technical Configuration Guide for the ERS 4500, 5000, BCM 50, 450, CS1000, CS2100 and SRG 50* (NN48500-576).

The Ethernet Routing Switch 8600 can be configured as a core switch only. There is no edge configuration for Nortel Automatic QoS on the Ethernet Routing Switch 8600. Presently, when used as a core switch for Nortel Automatic QoS with either the Ethernet Routing Switch 4500 or Ethernet Routing Switch 5000 as an edge switch, only Nortel Automatic QoS mixed mode is supported on the edge switch.

To configure Nortel Automatic QoS operation, configure the Nortel Voice Application with the proper Nortel Automatic QoS setting, enable DiffServ on the connected ingress port on the Ethernet Routing Switch, and then configure the port as a trusted core port. (The default operational value for Ethernet Routing Switch 8600 ports is core.)

### 802.1Q tagged packets

You cannot enable Nortel Automatic QoS for tagged packets on classic modules. Classic modules classify 802.1Q tagged ingress packets based on the 802.1p-bit values only, ignoring DSCP values. However, you can enable Nortel Automatic QoS on classic modules for untagged packets, because untagged packets have no p-bit values available and therefore use DSCP for classification.

To enable Nortel Automatic QoS support for tagged packets, you must equip the Ethernet Routing Switch 8600 with R or RS modules. These modules support an 802.1p-bit-override feature for tagged packets that allows the modules to ignore the 802.1p-bit and classify traffic based on the DSCP values instead. This feature is not supported on classic modules.

### Configuring Nortel Automatic QoS

To configure Nortel Automatic QoS, perform the following procedure.

#### Procedure steps

| Step | Action                                                                                                                                                                                                                                                                               |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>To enable Diffserv on a port, use the following CLI command:</p> <pre>config ethernet &lt;slot/port&gt; enable-diffserv true</pre> <p><b>OR</b></p> <p>use the following NNCLI command:</p> <pre>(config-if)# enable-diffserv [port &lt;portlist&gt;] enable</pre>                |
| 2    | <p>To enable a port as a trusted core port, use the following CLI command:</p> <pre>config ethernet &lt;slot/port&gt; access-diffserv false</pre> <p><b>OR</b></p> <p>use the following NNCLI command:</p> <pre>(config-if)# no access-diffserv [port &lt;portlist&gt;] enable</pre> |
| 3    | <p>For tagged ports on R/RS modules only, to enable 802.1p override, use the following CLI command:</p> <pre>config ethernet &lt;slot/port&gt; 802.1p-override enable</pre>                                                                                                          |

**OR**

use the following NNCLI command:

```
(config-if)# qos 802.1p-override enable
```

---

--End--

---

### **New operation for corrupt configuration loading**

In revisions of code prior to 5.0.1.0, when a switch boots and attempts to load a corrupt or invalid configuration file (with boot flag `verify-config true`; default value for the flag) and the switch does not have a valid backup configuration file, the switch boots with factory defaults. This means the switch runs as a pure L2 switch, with all ports active with default settings and all ports part of VLAN 1. This operation could cause improper network behavior, and a problematic switch without a valid configuration file loaded.

In the 5.1.0 release, a corrupt or invalid configuration file causes the switch to boot into 'monitor mode', in which no ports are active. The switch is isolated from the rest of the network so that the operation of the affected switch does not impact other switches and devices on the network. Any devices with a single attachment to this switch are still negatively impacted, but the impact is now limited to these devices. Devices that are dual attached, especially those SMLT/RSMLT attached, see little to no traffic interruption.

Access to an impacted switch running 5.1.0 code, for the purpose of repairing the switch, requires a console connection. This may require a local physical connection or a remote connection from a terminal server network that does not use the impacted switch.

Nortel recommends to set the `verify-config` boot flag to true, and to have backup configuration files available to prevent a switch from booting with factory defaults. Configuration of backup configuration files can be found in *Administration* (NN46205-605).

### **Core dumps**

Release 5.1 supports a core dump feature to provide Nortel with additional troubleshooting information in the rare instances where the switch/CPU resets itself.

If a crash should occur, a core dump file is now logged to an installed PCMCIA card containing sufficient memory to store the file. The name of the core dump log file is:

`core0Xdate_time.gz`, where X = to the SF/CPU slot number, and `date_time` is of format `yymmddtime`.

The system needs approximately 25MB of free-space for a non-SuperMezz system and 70MB of free-space for a SuperMezz system in order to be able to capture the core dump. A new log message is now added to the system to warn the user when the PCMCIA card memory falls below the threshold for the system to be able to fully write a core dump file.

You also have the ability to configure the writing of the core dump file to an external FTP server, to avoid concerns about having enough memory with the PCMCIA memory card. Since a TFTP client has file size limitations, FTP is the only option for externally saving core files.

### Configuring core dumps

You can enable or disable core dumps and configure the location where the core dump is saved. By default, core dumps are enabled and the file is saved to PCMCIA.

To enable core dumps with FTP, you must configure the **host user** and **host password** bootconfig parameters.

#### ATTENTION

If you configure the host user and password to nonnull values, all copying to and from the network (including image downloads) uses FTP instead of TFTP.

Because you can only configure one username and password combination for network copying and core dumps, be sure that the login for the remote system to which you are saving the core dumps is the same login used for the remote systems you are using for file copying.

| Step    | Action                                                                                                                                                                                                                                                                                          |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1       | <p>Configure the core dump using the following CLI command:</p> <pre>config bootconfig slot &lt;slots&gt; core-save {enable disable} [&lt;file&gt;]</pre> <p>OR</p> <p>using the following NNCLI command:</p> <pre>(config)# [no] boot config slot &lt;slots&gt; core-save [&lt;file&gt;]</pre> |
| --End-- |                                                                                                                                                                                                                                                                                                 |

**Nortel qualified DWDM XFPs**

Release 5.1 provides support for Nortel qualified DWDM XFPs. The following table provides wavelength information for each XFP. The frequency spacing is 100 GHz.

**ATTENTION**

The wavelength specifications provided by the XFP manufacturer and reported by the Ethernet Routing Switch 8600 system can vary slightly because of hexadecimal to decimal (and decimal to hexadecimal) rounding.

**Table 3**  
**DWDM XFP wavelengths**

| Part number | Wavelength and frequency            |
|-------------|-------------------------------------|
| NTK587AE-E5 | 1530.33 nm (195.90 Terahertz [THz]) |
| NTK587AG-E5 | 1531.12 nm (195.80 THz)             |
| NTK587AJ-E5 | 1531.90 nm (195.70 THz)             |
| NTK587AL-E5 | 1532.68 nm (195.60 THz)             |
| NTK587AN-E5 | 1533.47 nm (195.50 THz)             |
| NTK587AQ-E5 | 1534.25 nm (195.40 THz)             |
| NTK587AS-E5 | 1535.04 nm (195.30 THz)             |
| NTK587AU-E5 | 1535.82 nm (195.20 THz)             |
| NTK587AW-E5 | 1536.61 nm (195.10 THz)             |
| NTK587AY-E5 | 1537.40 nm (195.0 THz)              |
| NTK587BA-E5 | 1538.19 nm (194.9 THz)              |
| NTK587BC-E5 | 1538.98 nm (194.8 THz)              |
| NTK587BE-E5 | 1539.77 nm (194.7 THz)              |
| NTK587BG-E5 | 1540.56 nm (194.6 THz)              |
| NTK587BJ-E5 | 1541.35 nm (194.5 THz)              |
| NTK587BL-E5 | 1542.14 nm (194.4 THz)              |
| NTK587BN-E5 | 1542.94 nm (194.3 THz)              |
| NTK587BQ-E5 | 1543.73 nm (194.2 THz)              |
| NTK587BS-E5 | 1544.53 nm (194.1 THz)              |
| NTK587BU-E5 | 1545.32 nm (194.0 THz)              |

For more information, see *Nortel Ethernet Routing Switch 8600 Installation — SFP, XFP, GBIC, and OADM Hardware Components* (NN46205-320).


**Support for 100Base-FX SFP (AA1419074)**

The 5.1.0 release adds support for the 100Base-FX SFP for either the 8648GBRS or 8634XGRS module SFP ports. This support is just with these specific RS-modules.

## Support for new re-spin version of 8692SF/CPU (no part number or PEC change)

Release 5.0.1.0 and later provide support for a new re-spin version of the 8692SF/CPU. This re-spin need was created by the need to change parts associated with the PCMCIA controller slot. This change just affects this area of the 8692SF/CPU, and has no other form or functional impact, and therefore there is no PEC change. This affects PEC DS1404065-E5 and DS1404066-E5. The new re-spin version module can be identified by the PCMCIA card slot now having front panel artwork that shows "PC Card" versus older versions of the module show "PCMCIA".

The hardware was modified to address a part change associated with the PCMCIA controller. The new version of the hardware can be identified by the Hardware Rev 5.2 or by viewing the front panel text, shown below, which now shows "PC Card" for this new version vs. "PCMCIA" for the earlier versions.



The new version requires 5.0.1.0 code or higher, or 4.1.8.2 code or higher in the 4.1.x stream. All future code releases will also have support for this version of the module. Should you run this version of the module with lower numbered code releases, the module will boot and come up, but the PCMCIA/PC card slot will not function, and such a usage is not supported.

## Software licensing

The following table describes the license required to use specific features.

For more information about using the License Bank, see *Nortel Ethernet Routing Switch 8600 Administration* (NN46205-605).

**Table 4**  
License and features

| Base license                                                                                                                | Advanced license                                                                                                                                                                     | Premier license                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>Basic features</li> <li>IPV6 management</li> <li>BGPv4 for up to 10 peers</li> </ul> | <ul style="list-style-type: none"> <li>Border Gateway Protocol version 4 (BGPv4) for more than 10 Peers</li> <li>Bidirectional Forwarding Detection</li> <li>IPv6 Routing</li> </ul> | <ul style="list-style-type: none"> <li>all Advanced License features</li> <li>Virtual Routing and Forwarding, Lite version (VRF-Lite)</li> </ul> |

| Base license | Advanced license                                                                                                                         | Premier license                                                                                                                                                                                                                                                                                                                                      |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              | <ul style="list-style-type: none"> <li>• Multicast Source Discovery Protocol (MSDP)</li> <li>• Packet Capture function (PCAP)</li> </ul> | <ul style="list-style-type: none"> <li>• Multi-Protocol Border Gateway Protocol (MP-BGP)</li> <li>• IP-Virtual Private Network, Multi-Protocol Label Switching (RFC2547) (IP-VPN MPLS RFC2547)</li> <li>• IP-Virtual Private Network-Lite (IP-VPN-Lite – IP in IP)</li> <li>• Multicast virtualization for VRF-Lite (IGMP and PIM-SM/SSM)</li> </ul> |

Ethernet Routing Switch 8600 Release 5.1 and above includes a Premier trial license that is valid for 60 days from the date of install. After 60 days, the license expires and licensed features are no longer available after the system restarts.

## File names for this release

This section describes the Ethernet Routing Switch 8600 Software Release 5.1 software files.

Before you upgrade, verify the MD5 signature for each new file. For upgrade procedures, see *Nortel Ethernet Routing Switch 8600 Upgrades — Software Release 5.1* (NN46205-400).

**Table 5**  
**Release 5.1 software files**

| 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_5100.tar.gz | 59 M          |
| <b>Ethernet Routing Switch images</b> |                                                                                               |                  |               |
| Boot monitor image                    | CPU and switch fabric firmware                                                                | p80b5100.img     | 1.1 M         |
| Run-time image                        | Run-time image                                                                                | p80a5100.img     | 12 M          |
| Run-time image for R modules          | Image for R modules                                                                           | p80j5100.dld     | 1.4 M         |

**Table 5**  
**Release 5.1 software files (cont'd.)**

| Module or file type                                                     | Description                                                              | File name                                         | Size in bytes |
|-------------------------------------------------------------------------|--------------------------------------------------------------------------|---------------------------------------------------|---------------|
| Run-time image for RS modules                                           | Run-time image for RS modules                                            | p80k5100.dld                                      | 1.5 M         |
| Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz) | Image for the SuperMezz card                                             | p80m5100.img                                      | 12 M          |
| 3DES                                                                    | Encryption module for privacy protocol with Secure Shell (SSH)           | p80c5100.img                                      | 55 K          |
| AES                                                                     | Encryption module for privacy protocol for SNMPv3. Includes AES and 3DES | p80c5100.aes (this image includes the 3DES image) | 26 K          |
| MIB                                                                     | MIB files                                                                | p80a5100.mib                                      | 4.0 M         |
| MIB (zip file)                                                          | Zip file containing MIBs                                                 | p80a5100.mib.zip                                  | 659 K         |
| MD5 checksum file                                                       | md5 checksums of all Release 5.1 software files                          | p80a5100.md5                                      | 1.3 K         |
| Runtime image for ATM                                                   | Runtime image for the ATM module                                         | p80t5100.dld                                      | 885 K         |
| Runtime image for POS                                                   | Runtime image for the POS module                                         | p80p5100.dld                                      | 685 K         |
| <b>Firmware images</b>                                                  |                                                                          |                                                   |               |
| FOQ for R modules                                                       | Feedback output queueing FPGA firmware                                   | foq267.xsvf                                       | 5.1 M         |
| BMC for R modules                                                       | BAP memory controller FPGA firmware                                      | bmc776.xsvf                                       | 2.5 M         |
| DPC for R modules                                                       | Dual port controller FPGA firmware                                       | dpc184.xsvf                                       | 2.5 M         |
| PIM8630GBR                                                              | Programmable I/O module FPGA firmware; for the 8630GBR only              | PI_769.xsvf                                       | 2.2 M         |
| Firmware for RS modules                                                 | Contains FOQ, BMC, DPC, mirroring, and loopback images                   | rs_dpm_fpga.bin                                   | 4.3 M         |

**Table 5**  
**Release 5.1 software files (cont'd.)**

| Module or file type                                | Description                                                                              | File name                 | Size in bytes |
|----------------------------------------------------|------------------------------------------------------------------------------------------|---------------------------|---------------|
| PIM images for RS modules                          | PIM FPGA firmware required for 8612XLRS module only                                      | pim8612XLRS.bin           | 59 K          |
|                                                    | PIM FPGA firmware required for 8634XGRS module only                                      | pim8634XGRS.bin           | 76 K          |
|                                                    | PIM FPGA firmware required for 8648GBRS module only                                      | pim8648GBRS.bin           | 78 K          |
|                                                    | PIM FPGA firmware required for 8648GTRS module only                                      | pim8648GTRS.bin           | 53 K          |
| <b>SSL images</b>                                  |                                                                                          |                           |               |
| SSL cluster upgrade                                | Ethernet Routing Switch 8600 clustered SSL modules self-installing runtime image/upgrade | p80s5100.pkg              | 5.7 M         |
| SSL boot monitor                                   | Ethernet Routing Switch 8600 SSL module boot monitor                                     | p80s5100.img              | 7.2 M         |
| SSL upgrade instructions                           | Ethernet Routing Switch 8600 SSL upgrade instructions                                    | p80s5100.upgrade          | 1.4 K         |
| SSL installation instructions                      | Ethernet Routing Switch 8600 SSL installation instructions                               | p80s5100.install          | 2.8 K         |
| SSL diagnostics                                    | Ethernet Routing Switch 8600 SSL diagnostics                                             | p80s5100.diag             | 19 M          |
| <b>WSM images for Ethernet Routing Switch 8600</b> |                                                                                          |                           |               |
| WebOS firmware image                               | WSM WebOS v10.0.34.0 firmware image                                                      | wsm1003400_mp.img         | 826 K         |
| WebOS binary                                       | WSM WebOS v10.0.34.0 binary image                                                        | wsm1003400_bin.img        | 1.3 M         |
| WebOS boot image                                   | WSM WebOS v10.0.34.0 boot image                                                          | wsm1003400_boot.img       | 42 K          |
| <b>Device Manager images</b>                       |                                                                                          |                           |               |
| Solaris for SPARC image                            | Device Manager software image                                                            | jdm_6180_solaris_sparc.sh | 237 137 138   |

**Table 5**  
**Release 5.1 software files (cont'd.)**

| Module or file type                   | Description                                                                         | File name                                                                | Size in bytes |
|---------------------------------------|-------------------------------------------------------------------------------------|--------------------------------------------------------------------------|---------------|
| Microsoft Windows image               | Device Manager software image                                                       | jdm_6180.exe                                                             | 213 143 936   |
| Linux image                           | Device Manager software image                                                       | jdm_6180_linux.sh                                                        | 216 329 458   |
| <b>Service Delivery Module images</b> |                                                                                     |                                                                          |               |
| 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<br>NSF5100_2.3.7.0_SDM_R65.iso               |               |
|                                       | Upgrade package for the NSF Service Delivery Module Firewall                        | NSF5100_2.3.7.0_SDM_R60.pk<br>g.gz<br>NSF5100_2.3.7.0_SDM_R65.p<br>kg.gz |               |
| Service Delivery Module TPS           | Boot image for TPS Intrusion Sensor                                                 | Nortel_TPS_Intrusion_Sensor-S<br>DM-v4.5.0-627-Install.iso               |               |
|                                       | Boot ISO for TPS Defense Center booting from CD-ROM                                 | NortelTPSDefenseCenter_2x70v<br>4.5.0_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_<br>to_4.5.1_Upgrade-47.sh                   |               |
|                                       | IS upgrade download verification file.                                              | Nortel_TPS_IS_Upgrade_4.5.0_<br>to_4.5.1_Upgrade-47.sh.md5               |               |
| <b>Trace files</b>                    |                                                                                     |                                                                          |               |
| MPLS trace file                       | Trace file for MPLS. This is autogenerated and appears on the PCMCIA after upgrade. | nbpdtrc.lo0                                                              | variable      |

## Important information and restrictions

This section contains important information and restrictions that you should consider before you upgrade to Release 5.1.

### Important information and restrictions navigation

- [“SuperMezz, CPU memory, and upgrades” \(page 36\)](#)
- [“Nonsupported modules for Release 5.0 and later” \(page 39\)](#)
- [“I/O module considerations” \(page 39\)](#)
- [“Mixed-module MLT/LAG support for Release 5.0 and later” \(page 39\)](#)

- [“Fixes from previous releases” \(page 40\)](#)
- [“Supported upgrade paths” \(page 41\)](#)
- [“General upgrade considerations” \(page 41\)](#)
- [“Upgrade considerations for Release 5.x” \(page 41\)](#)
- [“Upgrade considerations: DOSFS with upgrades from pre-Release 5.0” \(page 41\)](#)
- [“Upgrade considerations: Power Management” \(page 42\)](#)
- [“Upgrade considerations: IST” \(page 45\)](#)
- [“Changes to options of license file installation procedures for Release 5.1” \(page 46\)](#)
- [“Considerations for upgrades from pre-5.0.1.0 release code” \(page 47\)](#)

### **SuperMezz, CPU memory, and upgrades**

If you want to use MPLS, VPN Lite, VPN, VRF Lite, or IPv6 features, you must install SuperMezz. The software prohibits configuration of these features unless a SuperMezz is installed.

You can upgrade to Release 5.0 or later without a SuperMezz installed, however, in some cases customers with large or complex configurations may encounter some scaling concerns. Using SuperMezz increases the available CPU processing power.

For Release 5.0 or later, Nortel recommends that the PCMCIA card in the PCMCIA slot be at least 256 MB.

Because of the new software features added for Release 5.0, such as VRF Lite, MPLS, and IP VPN, Release 5.0 and later require more CPU memory (RAM) than previous releases. In addition to this, a new extensive Core Dump functionality has been integrated and is enabled by default.

To determine memory usage, execute the **show sys perf** command before you upgrade. If the DramFree is below a predefined threshold, or other conditions are met, then SuperMezz is required before you upgrade to Release 5.0 or later. The DramFree threshold is 120 MB. This threshold is based on the following:

- 60 MB of DRAM is required for 5.0 or later software
- 50 to 60 MB of free DRAM is required for proper system operation (approximately 20 to 25% of 256 MB)

Certain network designs, environments, and system configurations can be described as “highly-scaled”. In such circumstances there can be a combination of resource requirements - in terms of CPU Memory and CPU Processing - and configuration demands that make using the SuperMezz highly recommended if not mandatory. You should examine all aspects of your installation to determine if it is a “highly-scaled” environment.


See the following figures to determine if SuperMezz is required.

The following figure shows the decision-making tree for E or M or mixed-module systems. A mixed-module system has both E or M modules and R or RS modules.

### ATTENTION


In the following figures, the DramFree threshold is 120 MB.

**Figure 1**  
**Decision-making tree for E or M or mixed-module systems**



The following figure shows the decision-making tree for R and RS module systems.

**Figure 2**  
**Decision-making tree for R and RS module systems**



Disabling Core Dump can free 16 to 20 MB of DRAM. Nortel recommends that you do not disable Core Dump unless absolutely necessary. Core Dump provides valuable troubleshooting and debugging information. Disabling Core Dump requires a switch reboot.

To disable Core Dump:

- from running configuration CLI mode:  
`config bootconfig slot <slot-num> core-save disable`  
 where <slot-num> is the slot number of the SF/CPU module, either 5 or 6.
- from boot monitor CLI mode:  
`slot <slot-num> core-save disable`
- from running configuration NNCLI mode:

```
no boot config slot <5,6> core-save
```

- from boot monitor NNCLI mode:  

```
no slot <5,6> core-save
```

### **Nonsupported modules for Release 5.0 and later**

Release 5.0 and later does not support the 1-port 10 Gigabit Ethernet 8681XLW and 8681XLR modules, the 8690 SF/CPU modules, or any pre-E modules.

### **I/O module considerations**

The 8648GTR module does not support packets larger than 9188 bytes at 100 Mbps. At 1000 Mbps, frames larger than 9188 bytes are supported.

R and RS modules require the 8692SF

### **Mixed-module MLT/LAG support for Release 5.0 and later**

Mixed-module multilink trunks that consist of both R and RS module and Classic module (E and M module) ports is fully supported. However, mixed-module multilink trunks are only fully supported for common features that function for both types of modules.

Many parameters must be set at the port level within the multilink trunk. VLAN assignment, tagging (802.1d) status, and the Spanning Tree Protocol are examples of some parameters that do not require configuration at the port level, but instead are configured at the multilink trunk level. If a parameter is only available for a subset of the multilink trunk ports, then the use of this parameter is not supported. With such use, inconsistent behavior may occur.

Speed and duplex settings are examples of port-level parameters that are supported for all modules, and therefore can always be used. However, common port-level parameters must be identically configured for all ports within the multilink trunk.

Multicast and broadcast rate limiting is available on both R and Classic modules, but its operation depends upon the module type. Therefore, use of rate limiting is not recommended for mixed-module multilink trunks; its use may lead to inconsistent or unknown behavior.

Egress queue size configuration is only available with R and RS modules and therefore cannot be used within a mixed-module multilink trunk. The egress queue size cannot be applied to Classic module ports.

Apply caution when using filters within a mixed-module multilink trunk. Filter operation is different for different module types. If filters must be used, ensure that they apply the same rules and actions to both Classic

and R series module port types. Because the number of filter permutations is extensive, there is little or no software control or checking. Therefore, it is the responsibility of the user to properly configure any mixed-module multilink trunks. Any questions or concerns for such configurations should be addressed to Nortel prior to implementation.

MLT hashing is a function of the individual ingress port whether the port is a single port or part of a multilink trunk (either same or mixed-module configuration). Therefore, how the packet is hashed is completely independent of the type of egress multilink trunk.

With this release, you can use different module types for MLT/LAG (Link Aggregation Group) configurations. However, ensure that all ports in the MLT/LAG use the same port speed, for example, 1 Gbps, even if Auto-Negotiation is used. The software does not enforce this requirement. Nortel recommends that you use CANA to ensure proper speed negotiation in mixed-port type scenarios.

To maintain LAG stability during failover, use CANA: configure the advertised speed to be the same for all LACP links. For 10/100/1000 Mbps ports, ensure that CANA uses one particular setting, for example, 1000-full or 100-full. Otherwise, a remote device could restart Auto-Negotiation and the link could use a different capability.

It is important that each port uses only one speed and duplex mode. This way, all links in Up state are guaranteed to have the same capabilities. If Auto-Negotiation and CANA are not used, the same speed and duplex mode settings should be used on all ports of the MLT. This is true for both 10/100/1000 Mbps modules, and for 10/100 Mbps Classic modules that do not support CANA.

In LACP LAGs, if one of the links goes operationally down (for example, because of speed mismatch, cable removal, Auto-Negotiation problems, or other failures), the link remains in the LAG (that is, in the MLT). Accordingly, show commands show the port as part of the MLT. In LACP LAGs, if the link is up, and the Auto-Negotiation, speed, or duplex settings do not match, then that link is removed from the LAG, or the LAG fails.

### **Fixes from previous releases**

The Ethernet Routing Switch 8600 Software Release 5.1 incorporates fixes from releases up to and including 4.1.8 and 5.0.1.

## Supported upgrade paths

The Ethernet Routing Switch 8600 Software Release 5.1 supports upgrades from the following earlier releases:

- 4.1.x
- 5.0.x

If you want to upgrade to 5.1 from releases prior to 4.1.x, first upgrade to 4.1.x or 5.0.x and then upgrade to 5.1.

## General upgrade considerations

The configuration file generated with Ethernet Routing Switch 8600 Software Release 5.0 and later contains options that are not backward-compatible with Ethernet Routing Switch 8600 Software Releases 3.0.x, 3.1.x, 3.2.x, 3.3.x, 3.5.x, 3.7.x or 4.0.x.

Loading a Release 5.0 or later configuration file on a 3.0.x, 3.1.x, 3.2.x, or 3.3.x runtime image generates errors and causes the image to stop loading the configuration file. The system will load with a default configuration.

Downgrades require previously saved configuration files (boot.cfg and config.cfg) and may require the removal of R and RS modules prior to downgrade.

## Upgrade considerations for Release 5.x

Before you upgrade, read *Nortel Ethernet Routing Switch 8600 Upgrades — Software Release 5.1* (NN46205-400) and follow the outlined procedures.

The first 5.x upgrade you perform requires you to reformat the DOSFS for the PCMCIA and flash. Steps are included in the upgrade procedures. See [“Upgrade considerations: DOSFS with upgrades from pre-Release 5.0”](#) (page 41).

You must take into consideration Power Management for this release; for more information, see [“Upgrade considerations: Power Management”](#) (page 42).

## Upgrade considerations: DOSFS with upgrades from pre-Release 5.0

Release 5.0 introduces a unique signature to the Disk Operating System File System (DOSFS) volume label generated during `dos-format` and `format-flash` operations. This label provides clear identification about which DOSFS devices have been formatted with the latest DOSFS source code.

When you upgrade from pre-Release 5.0 software and boot an image with Release 5.0 or later, you may see boot messages like:

The /flash device mounted successfully, but it appears to have been formatted with pre-Release 5.0 file system code. Nortel recommends backing up the files from /flash, and executing `dos-format /flash` to bring the file system on the /flash device to the latest ERS8600 baseline.

If you receive this message, Nortel recommends that you perform a one-time reformat of the DOSFS device (using **`dos-format`**) to set the DOSFS baseline. This is part of the upgrade procedures.

The one-time DOS reformat erases all files on the DOSFS device. Nortel recommends that you back up all files from the DOSFS device, reformat the device, and replace all files.

### **Upgrade considerations: Power Management**

The addition of the Power Management feature to Release 5.0 may require you to take special steps before you upgrade.

When you upgrade to Release 5.0 or later, Power Management is enabled by default. If Power Management detects that there are not enough power supplies in the system to successfully run the system, it shuts down the lowest-priority modules. This does not occur if you have enough available power.

You can calculate the number of power supplies required for your Ethernet Routing Switch 8600 system. To determine the number of power supplies required for your switch configuration, use the *Nortel ERS8600 Power Supply Calculator* (NN48500-519). This is available on the Nortel Technical Support Web site at [www.nortel.com/documentation](http://www.nortel.com/documentation). Choose Routers & Routing Switches, and then Ethernet Routing Switch 8600. In the DOCUMENTATION, Operations section, click the Configuration link. Navigate through the list until you find the calculator.

For Power Management configuration and conceptual information, see *Ethernet Routing Switch 8600 Administration* (NN46205-605).

### **Power Management operations**

With Power Management, when the switch boots, users are notified if there is redundant power available in the system. This notification is based on the available power provided by the power supplies as compared to the power requirements of the installed modules.

R and RS modules are not brought up if there is insufficient power available. Although there is an override capability available, this should only be used for short periods of time or in emergencies—operating a chassis in an underpowered condition can lead to unpredictable results. Power management does not apply to switch fabric modules (8691SF/CPU or 8692SF/CPU), and does not apply to prior generations of I/O modules (pre-E, E, and M).

The amount of system power is calculated based on the number, type, and input source voltage of the power supplies in the chassis. For 8005AC or 8005DI AC supplies, the system detects whether the supply is sourced with 110 V or 220 V and uses the corresponding output power. For all other power supplies, the power output is the same, regardless of source voltage.

Switch fabrics are allotted highest priority and always power up. Pre-E, E, and M modules always power up. R and RS modules power up if there is sufficient power remaining to do so. If there is insufficient power to bring all R and RS modules online, they are powered up based on slot priority. By default, R and RS modules are powered up starting at slot 1 until there is insufficient power to bring the next module online.

If a chassis boots up and there are modules that are not online due to insufficient power, adding an additional power supply does not bring the modules online automatically. To bring the modules online, the system must be rebooted, or the module must be removed and reinserted into the chassis.

If a system boots and power supply failure occurs, one of the two following conditions result:

1. A system with redundant power continues to operate normally. The redundant power configuration compensates for a power supply failure.
2. A system with no redundant power continues to operate, however, if there is insufficient power to support all modules, an SNMP trap and syslog message are sent every five minutes notifying the user that the system is operating in an underpowered condition. This situation must be corrected as soon as possible.

### **Disabling power and cooling management**

Disable Power Management to successfully upgrade even though not enough power supplies are installed to run all I/O modules.

If you already have enough power supplies, you do not need to disable Power Management.

You can calculate the number of power supplies required for your Ethernet Routing Switch 8600 system. To determine the number of power supplies required for your switch configuration, use the *Power Supply Calculator for ERS 8600* (NN48500-519). This is available on the Nortel Technical Support Web site at [www.nortel.com/documentation](http://www.nortel.com/documentation). Choose Routers & Routing Switches, and then Ethernet Routing Switch 8600.

**ATTENTION**

Nortel recommends that you do not disable Power Management, and that you instead install the required power supplies before upgrade. However, if you must disable Power Management for a short period of time, install the required supplies as quickly as possible.

RS modules do not come up when the cooling module is not installed.

**ATTENTION**

Although you can override the fan check for the cooling module, this should only be done for short periods of time or in emergencies—operating a chassis with RS modules without the cooling module can lead to unpredictable results.

| Step                           | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                              | Save the pre-5.1 configuration file.<br><br><code>save config.cfg</code>                                                                                                                                                                                                                                                                                                                                                                                                         |
| 2                              | Edit the configuration file offline using an editor like VI or EMACS. You can either: <ul style="list-style-type: none"><li>• Use the CLI to edit the file on the switch (the switch has a built-in VI-like editor). Use the <code>edit config.cfg</code> command.</li><li>• Save the file as an ASCII file and transfer to another device for editing with a text editor like Notepad.</li><li>• Transfer the file to a device and edit with VI or EMACS-like editor.</li></ul> |
| 3                              | In the configuration file, add the following lines to the end of the flags section:<br><br><code>#!power power-check-enable false</code><br><br><code>#!power fan-check-enable false</code><br><br>See the following job aid for an example of correct placement of these commands.                                                                                                                                                                                              |
| 4                              | Save the file and, if you edited it off-switch, transfer the file back to the switch to use in the upgrade.                                                                                                                                                                                                                                                                                                                                                                      |
| <hr/> <div>--End--</div> <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Job aid: configuration file and command placement**

```
#
# MON MAY 19 22:43:41 2008 UTC
# box type : ERS-8010
# software version : REL5.0.0.0_B006
# monitor version : 5.0.0.0/006
# cli mode : 8600 CLI
#
#
# Asic Info :
# SlotNum|Name |CardType |MdaType |Parts Description
#
# Slot 1 -- 0x00000001 0x00000000
# Slot 2 -- 0x00000001 0x00000000
# Slot 3 -- 0x00000001 0x00000000
# Slot 4 8630GBR 0x2432511e 0x00000000 RSP=25 CLUE=2 F2I=1 F2E=1
FTMUX=17 CC=3 FOQ=266 DPC=184 BMC=776 PIM=257 MAC=4
# Slot 5 8692SF 0x200e0100 0x00000000 CPU: CPLD=19 MEZZ=4 SFM:
OP=3 TMUX=2 SWIP=23 FAD=16 CF=56
# Slot 6 -- 0x00000001 0x00000000
# Slot 7 -- 0x00000001 0x00000000
# Slot 8 -- 0x00000001 0x00000000
# Slot 9 -- 0x00000001 0x00000000
# Slot 10 -- 0x00000001 0x00000000
#
#!flags m-mode false
#!flags enhanced-operational-mode false
#!flags vlan-optimization-mode false
#!flags global-filter-ordering false
#!flags r-mode false
#!resource-reservation max-vlan false
#!resource-reservation multicast 2048
#!flags multicast-check-packet true
#!flags system-monitor true
#!record-reservation filter 4096
#!record-reservation ipmc 500
#!record-reservation local 2000
#!record-reservation mac 2000
#!record-reservation static-route 200
#!record-reservation vrrp 500
#!system-monitor monitoring-enable true
#!system-monitor detection-time 30
#!power power-check-enable false <----- ADD THIS LINE
#!power fan-check-enable false <----- ADD THIS LINE
```

**Upgrade considerations: IST**

After an IST peer is upgraded and restarted, wait until the entire system is stable prior to upgrading the other IST peer. Stabilization time depends on the complexity and size of the network (for example, the number of MAC and ARP records, routes, and the protocols used). Wait for the Layer 3

protocols, especially multicast protocols, to settle before you restart the other peer. If Layer 3 protocols are not in use, wait until the FDB and ARP tables on both peers report a similar number of entries.

## Changes to options of license file installation procedures for Release 5.1

After you upgrade to release 5.1, you must specify the name and location of your license file in the boot configuration file. If you do not specify the name and location of your license file, you can encounter issues with your licensed features.

### Procedure steps

| Step    | Action                                                                                                                                                                                                                                                                                   |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1       | <p>To specify the license file location, enter the following CLI command:</p> <pre>config bootconfig choice primary license-file &lt;file&gt;</pre> <p><b>OR</b></p> <p>enter the following NNCLI command:</p> <pre>(config)# boot config choice primary license-file &lt;file&gt;</pre> |
| <hr/>   |                                                                                                                                                                                                                                                                                          |
| --End-- |                                                                                                                                                                                                                                                                                          |

**Note:** The variable '<file>' supports the following values for the source of a license file on an ERS8600 switch:

- `/flash/<file_name>`
- `/pcmcia/<file_name>`
- `<a.b.c.d>:<file_name>`

## Considerations for upgrades from pre-5.0.1.0 release code

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.

**Note:** For switch cluster systems running 5.0.0.x code, an intermediate upgrade to 5.0.0.2 is required, versus a direct upgrade to 5.1.0. If not performed, direct console access will be required to recover the 'peer' switch cluster system still running 5.0.0.x code, after the first switch is upgraded. Direct upgrades are supported from all 4.1.x code, including 4.1.8.2 (CR Q01977204). See *Nortel Ethernet Routing Switch 8600 Upgrades — Software Release 5.0* (NN46205-400) for general steps to upgrade to 5.0.0.2.

When upgrading a pair of IST Core peer switches, special care must be taken due to SMLT architectural change. Unlike previous upgrades, when moving from pre-5.0.1.0 code to any 5.1.0 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-5.0.1.0 image. The user should now wait approximately 5 minutes until the following message is seen on the upgraded 5.1.0 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.
```

The second IST Core peer switch can then be upgraded. When the first switch, now running 5.1.0, sees its IST MLT connection go down, it unlocks all its ports and starts learning and then forwarding traffic. This upgrade process will cause some user traffic impact, even for the SMLT connected portion of the network. The second IST Core switch, after upgrading, should show the following messages in its log. These should be checked 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 the network is running 'normally', it is now recommended to power-cycle (reboot only for single SSF/CPU systems) the one peer switch which was upgraded first so that a complete sync occurs.

After upgrading to 5.1.0 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.

**ATTENTION**

In the upgrade procedures, boot the Secondary SF/CPU first and then wait 10 seconds and then immediately boot the Master SF/CPU. Do NOT wait for the Secondary SF/CPU to come up completely before booting the Master SF/CPU, or the Secondary SF/CPU may hang.

This new upgrade operation affects SMLT based networks moving from pre-5.0.1.0 code to 5.1.0 or higher code. If moving from 4.1.8.x to 5.0 based code, it is recommended to move to 5.1.0 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.

Upgrades from 4.1.x to 5.1.0 are 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 entries are returned, please contact Nortel Support and request Level 2 GNPS assistance.



**CAUTION**

If after upgrading to 5.1.0 code, the user wants to downgrade to some other 5.0.x code, that downgrade MUST be done to 5.0.0.2 code instead of 5.0.0.0 or 5.0.0.1 code. If present, the user may still use the same 5.0.0.0 or 5.0.0.1 config files when performing the downgrade but the downgrade code release must be 5.0.0.2. Downgrades to 4.1.x code are as normal supported. (Q01971801)

**CAUTION**

Users using BGP need to be aware of the following limitations regarding upgrading to 5.1.0 code release; this limitation also applied to any previous 5.0.x release, but was unknown previously. For any user using the add-as-path command in 4.x or earlier releases, a direct upgrade to 5.x code (or specifically 5.1.0 code) will create issues with your BGP operation, as the format for this command has changed in 5.0.x and all future code releases. This usage of this command can be confirmed by looking at your current 4.x based configuration file (config.cfg by default) by using either CLI command show config or more /flash/config.cfg, and looking for entries under:

```
# IP AS LIST CONFIGURATION #
```

Entries such as this indicate usage of the command:

```
ip as-list 1 create ip as-list 1 add-as-path 100
permit "64521"
```

With 5.0.x code, the two commands have been replaced by a single command of format:

```
ip as-list <as-list id; 1-1024> create <member id
in as-path; 0-65535> permit "<as-path: 0-65535>"
```

Prior to upgrading to 5.x code (or later), if such config entries are in a 4.x config file, those entries must be manually converted to 5.x format before upgrading; the upgrade to 5.x code does not convert this command structure properly. Since both the 4.x and 5.x code files are plain ASCII text, the 4.x config file can be copied to any text editor (or edited locally on the 8600 switch with its Unix VI editor), edited (for example with M/S Word) and then copied back before upgrading. The above 4.x config example:

```
ip as-list 1 create ip as-list 1 add-as-path 100
permit "64521"
```

(Q01977204)

Release 5.1.0 has been verified to support the ability to migrate from configuration files created with any software release from 4.1.x and above.

Downgrades require previously saved configuration files (boot.cfg and config.cfg) and may require the removal of R and RS modules prior to downgrade.

### Switch cluster upgrade considerations for pre-5.0.1.0 release code

Due to the SMLT architectural changes introduced with release 5.0.1 and described in the previous section, you must follow specific upgrade steps in some scenarios when upgrading SMLT switch clusters from pre-release 5.0.1 to release 5.1 and later.


For device management during an upgrade, you can use one of the following options:

1. Direct serial console connection to the switch
2. Telnet access to the management IP
3. Telnet access to any of the in-band IP addresses on the switch

In scenarios 1 and 2, you can manage the switch effectively at all times during the upgrade, and therefore these scenarios require no additional considerations. However, in scenario 3, due to the SMLT architectural changes, you can lose Telnet connectivity during the upgrade of the IST peers unless you follow the proper steps.

Consider the following figure, showing a triangle SMLT setup. In this case, the user intends to upgrade the IST peers (that are currently running 5.0.0.2) to 5.1 or later.


**Figure 3**  
**SMLT upgrade scenario**



Assume the user Telnets from the PC to manage switch A and switch B. When the Telnet traffic generated by the PC arrives at switch C, depending on the MLT hashing algorithm, the traffic can be hashed to the link toward switch A or switch B. So, it is possible to have a situation where the Telnet management traffic destined for switch A flows through switch B and vice-versa.

Assume that the user upgrades switch A to 5.1. Due to the new SMLT behavior, the network diagram now looks like the following figure.

**Figure 4**  
**SMLT upgrade scenario after upgrading switch A to 5.1**



In this situation the PC cannot communicate with switch A, and as a result Telnet access to switch A is unavailable. For in-band management, you can alternatively Telnet first into switch B, and then Telnet to switch A from there.

The following are the recommended steps to perform this upgrade procedure while using Telnet in-band management:

1. Telnet to switch B from the PC
2. From switch B, Telnet to switch A
3. Upgrade switch A to 5.1.0, following the normal upgrade process. At this point, your Telnet session to switch A is lost, and eventually times out. After approximately a minute, Telnet to switch A again. This allows you to check the log messages on switch A. (At this point, you can possibly lose the Telnet connectivity to B in some situations depending on the MLT hashing occurring on switch C. If this occurs, re-open a Telnet connection to switch B.)
4. Upgrade switch B to 5.1 following the normal upgrade process. At this point, your Telnet session to switch B is lost. You can open a new Telnet session to switch A. After switch B completes the upgrade, you can then establish connectivity with switch B, either via Telnet from switch A, or via Telnet from the PC.

The same procedure applies for warm standby and hot standby scenarios. You must follow the upgrade directions for warm and hot standby cases provided in the upgrade document for individual chassis.

**Note:** If switch A and switch B are running 5.0.0.x (where x is less than 2), the switches MUST be upgraded to 5.0.0.2 before upgrading to 5.1.0.0 or later. This rule does not apply to upgrades to 5.1.0.0 from 5.0.1.0.

## Supported software and hardware scaling capabilities

The following table lists supported software and hardware scaling capabilities in Ethernet Routing Switch 8600 Software Release 5.1. The information in this table supersedes information contained in *Nortel Ethernet Routing Switch 8600 Planning and Engineering — Network Design* (NN46205-200), or any other document in the suite.

The capabilities described in this table are supported as individual protocols, not mixtures of protocols.

Nortel supports 25 Spanning Tree Groups (STG) in this release. Although you can configure up to 64 STGs (only 63 when a Web Switching Module is present), configurations including more than 25 STGs are not supported. If you need to configure more than 25 STGs, contact your Nortel Customer Support representative for more information about the support of this feature.

MLT is statically compliant with the 802.3ad standard (no support of LACP).

### ATTENTION

To install full Internet Routes, SuperMezz is required. Without SuperMezz, only 100 000 routes are supported.

**Table 6**  
**Supported scaling capabilities**

|                           | Maximum number supported 8692SF without SuperMezz (R or RS series modules) | Maximum number supported 8692SF with SuperMezz (R or RS series modules) | Maximum number supported (E and M modules) |
|---------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|--------------------------------------------|
| <i>Layer 2</i>            |                                                                            |                                                                         |                                            |
| MAC address table entries | 64 000<br>32 000 when SMLT is used                                         | 64 000<br>32 000 when SMLT is used                                      | E modules: 25 000<br>M modules: 119 000    |

**Table 6**  
**Supported scaling capabilities (cont'd.)**

|                                                                                     | Maximum number supported 8692SF without SuperMezz (R or RS series modules) | Maximum number supported 8692SF with SuperMezz (R or RS series modules) | Maximum number supported (E and M modules)                              |
|-------------------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------|
| VLANs (port-protocol-, and IEEE 802.1Q-based)                                       | 4000 when max VLAN feature enabled                                         | 4000                                                                    | 1972                                                                    |
| IP subnet-based VLANs                                                               | 800                                                                        | 800                                                                     | E modules: 200<br>M modules: 800                                        |
| Ports per Link Aggregation Group (LAG, MLT)                                         | 8                                                                          | 8                                                                       | 8                                                                       |
| Aggregation groups<br>802.3ad aggregation groups<br>Multi Link Trunking (MLT) group | NonR mode: 32<br>R mode: 128                                               | NonR mode: 32<br>R mode: 128                                            | 32                                                                      |
| SMLT links                                                                          | R mode: 128                                                                | R mode: 128                                                             | 32                                                                      |
| SLT (single link SMLT)                                                              | 382                                                                        | 382                                                                     | 382                                                                     |
| VLANs on SMLT/IST link                                                              | R mode with Max VLAN feature enabled: 2000                                 | R mode with Max VLAN feature enabled: 2000                              | 990 (with Enhanced Operational mode enabled)                            |
| RSMLT per VLAN                                                                      | 32 SMLT links with RSMLT-enabled VLANs                                     | 32 SMLT links with RSMLT-enabled VLANs                                  | 32 SMLT links with RSMLT-enabled VLANs                                  |
| RSTP/MSTP (number of ports)                                                         | 384, with 224 active. Configure the remaining interfaces with Edge mode    | 384, with 224 active. Configure the remaining interfaces with Edge mode | 384, with 224 active. Configure the remaining interfaces with Edge mode |
| MSTP instances                                                                      | 32                                                                         | 32                                                                      | 32                                                                      |
| <i>Advanced Filters</i>                                                             |                                                                            |                                                                         |                                                                         |
| ACLs for each system                                                                | 4000                                                                       | 4000                                                                    | N/A                                                                     |
| ACEs for each system                                                                | 1000                                                                       | 1000                                                                    | N/A                                                                     |
| ACEs for each ACL                                                                   | 1000                                                                       | 1000                                                                    | N/A                                                                     |
| ACEs for each port                                                                  | 2000:<br>500 inPort<br>500 inVLAN<br>500 outPort<br>500 outVLAN            | 2000:<br>500 inPort<br>500 inVLAN<br>500 outPort<br>500 outVLAN         | N/A                                                                     |

**Table 6**  
**Supported scaling capabilities (cont'd.)**

|                                                                                                      | Maximum number supported 8692SF without SuperMezz (R or RS series modules) | Maximum number supported 8692SF with SuperMezz (R or RS series modules) | Maximum number supported (E and M modules)                                                 |
|------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| <i>IP, IP VPN/MPLS, IP VPN Lite, VRF Lite</i>                                                        |                                                                            |                                                                         |                                                                                            |
| IP interfaces (VLAN- and brouter-based)                                                              | 1972                                                                       | 1972                                                                    | 1972                                                                                       |
| VRF instances                                                                                        | N/A                                                                        | 255                                                                     | N/A                                                                                        |
| ECMP routes                                                                                          | 5000                                                                       | 5000                                                                    | 5000                                                                                       |
| VRRP interfaces                                                                                      | 255                                                                        | 255                                                                     | 255                                                                                        |
| IP forwarding table (Hardware)                                                                       | 120 000                                                                    | 250 000                                                                 | E modules: 20 000 forwarding routes<br>M modules: 119 000 forwarding routes                |
| BGP/mBGP peers                                                                                       | 10                                                                         | 250                                                                     | 10                                                                                         |
| iBGP instances                                                                                       | on GRT                                                                     | on GRT                                                                  | on GRT                                                                                     |
| eBGP instances                                                                                       | on GRT                                                                     | on 256 VRFs (including GRT)                                             | on GRT                                                                                     |
| BGP forwarding routes<br>BGP routing information base (RIB)<br>BGP forwarding information base (FIB) | BGP FIB 120 000<br>BGP RIB 250 000                                         | BGP FIB 250 000<br>BGP RIB 500 000                                      | E modules: BGP FIB 20 000; BGP RIB 250 000;<br>M modules: BGP FIB 119 000, BGP RIB 250 000 |
| IP VPN routes (total routes for each system)                                                         | N/A                                                                        | 180 000                                                                 | N/A                                                                                        |
| IP VPN VRF instances                                                                                 | N/A                                                                        | 255                                                                     | N/A                                                                                        |
| Static ARP entries                                                                                   | 2048                                                                       | 2048 per VRF<br>10 000 per system                                       | 2048                                                                                       |
| Dynamic ARP entries                                                                                  | 32 000 in R mode                                                           | 32 000 in R mode                                                        | 16 000                                                                                     |
| DHCP relay instances (total for all VRFs)                                                            | 512                                                                        | 512                                                                     | 512                                                                                        |
| Static route entries                                                                                 | 2000                                                                       | 2000 per VRF<br>10 000 per system                                       | 2000                                                                                       |
| OSPF instances for each switch                                                                       | on GRT                                                                     | on 64 VRFs (including GRT)                                              | on GRT                                                                                     |
| OSPF areas for each switch                                                                           | 5                                                                          | 5 per VRF<br>24 per system                                              | 5                                                                                          |

**Table 6**  
**Supported scaling capabilities (cont'd.)**

|                                      | Maximum number supported 8692SF without SuperMezz (R or RS series modules) | Maximum number supported 8692SF with SuperMezz (R or RS series modules) | Maximum number supported (E and M modules) |
|--------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|--------------------------------------------|
| OSPF adjacencies for each switch     | 80                                                                         | 80<br>200 per system                                                    | 80                                         |
| OSPF routes                          | 20 000                                                                     | 20 000 per VRF<br>50 000 per system                                     | E modules 15 000<br>M modules 20 000       |
| OSPF interfaces                      | 238                                                                        | 500<br>500 per system                                                   | 238                                        |
| OSPF LSA packet maximum size         | 3000 bytes                                                                 | 6000 bytes                                                              | 3000 bytes                                 |
| RIP instances                        | on GRT                                                                     | on 64 VRFs (including GRT)                                              | on GRT                                     |
| RIP interfaces                       | 200                                                                        | 200                                                                     | 200                                        |
| RIP routes                           | 2500                                                                       | 2500 per VRF<br>10 000 per system                                       | 2500                                       |
| <i>Multiprotocol Label Switching</i> |                                                                            |                                                                         |                                            |
| MPLS LDP sessions                    | N/A                                                                        | 200                                                                     | N/A                                        |
| MPLS LDP LSPs                        | N/A                                                                        | 16 000                                                                  | N/A                                        |
| MPLS RSVP static LSPs                | 200                                                                        | 200                                                                     | N/A                                        |
| Tunnels                              | 2500                                                                       | 2500                                                                    | N/A                                        |
| <i>IP Multicast</i>                  |                                                                            |                                                                         |                                            |
| DVMRP passive interfaces             | 1200                                                                       | 1200                                                                    | 1200                                       |
| DVMRP active interfaces/neighbors    | 80                                                                         | 80                                                                      | 80                                         |
| DVMRP routes                         | 2500                                                                       | 2500                                                                    | 2500                                       |
| PIM instances                        | on GRT                                                                     | on 64 VRFs (including GRT)                                              | on GRT                                     |
| PIM active interfaces                | 200                                                                        | 200 (200 for all VRFs)                                                  | 200                                        |
| PIM passive interfaces               | 1972                                                                       | 1972 (2000 for all VRFs)                                                | 1972                                       |
| PIM neighbors                        | 80                                                                         | 80 (200 for all VRFs)                                                   | 80                                         |
| MSDP peers                           | 20                                                                         | 20                                                                      | 20                                         |

**Table 6**  
**Supported scaling capabilities (cont'd.)**

|                                                    | Maximum number supported 8692SF without SuperMezz (R or RS series modules) | Maximum number supported 8692SF with SuperMezz (R or RS series modules) | Maximum number supported (E and M modules) |
|----------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|--------------------------------------------|
| MSDP maximum SA messages                           | 6144                                                                       | 6144                                                                    | 6144                                       |
| Multicast streams: with SMLT/ without SMLT         | 500/1500                                                                   | 2000/4000                                                               | 500/1500                                   |
| <i>IPX</i>                                         |                                                                            |                                                                         |                                            |
| IPX interfaces                                     | N/A                                                                        | N/A                                                                     | 100                                        |
| IPX RIP routes                                     | N/A                                                                        | N/A                                                                     | 5000                                       |
| IPX SAP entries                                    | N/A                                                                        | N/A                                                                     | 7500                                       |
| <i>IPv6</i>                                        |                                                                            |                                                                         |                                            |
| IPv6 interfaces                                    | N/A                                                                        | 250                                                                     | N/A                                        |
| IPv6 tunnels                                       | N/A                                                                        | 350                                                                     | N/A                                        |
| IPv6 static routes                                 | N/A                                                                        | 2000                                                                    | N/A                                        |
| OSPFv3 areas                                       | N/A                                                                        | 5                                                                       | N/A                                        |
| OSPFv3 adjacencies                                 | N/A                                                                        | 80                                                                      | N/A                                        |
| OSPFv3 routes                                      | N/A                                                                        | 5000                                                                    | N/A                                        |
| <i>Operations, Administration, and Maintenance</i> |                                                                            |                                                                         |                                            |
| IPFIX                                              | 384 000 flows per chassis                                                  | 384 000 flows per chassis                                               | N/A                                        |
| RMON alarms with 4000K memory                      | 2630                                                                       | 2630                                                                    | 2630                                       |
| RMON events with 250K memory                       | 324                                                                        | 324                                                                     | 324                                        |
| RMON events with 4000K memory                      | 5206                                                                       | 5206                                                                    | 5206                                       |
| RMON Ethernet statistics with 250K memory          | 230                                                                        | 230                                                                     | 230                                        |
| RMON Ethernet statistics with 4000K memory         | 4590                                                                       | 4590                                                                    | 4590                                       |

**CAUTION****Risk of traffic loss**

SuperMezz is required for configurations that carry over 4000 multicast routes.

## Hardware and software compatibility

The following table describes your hardware and the minimum Ethernet Routing Switch 8600 software version required to support the hardware.

**Table 7**  
**Chassis, power supply, and SF/CPU compatibility**

| Item                        |                                                   |                                                                                                    | Minimum software version | Part number                    |
|-----------------------------|---------------------------------------------------|----------------------------------------------------------------------------------------------------|--------------------------|--------------------------------|
| <b>Chassis</b>              |                                                   |                                                                                                    |                          |                                |
|                             | 8010co                                            | 10-slot                                                                                            | 3.1.2                    | DS1402004-E5<br>DS1402004-E5GS |
|                             | 8010                                              | 10-slot                                                                                            | 3.0.0                    | DS1402001-E5<br>DS1402001-E5GS |
|                             | 8006                                              | 6-slot                                                                                             | 3.0.0                    | DS1402002-E5<br>DS1402002-E5GS |
| <b>Switching fabric/CPU</b> |                                                   |                                                                                                    |                          |                                |
|                             | 8691SF/CPU                                        | Switching fabric                                                                                   | 3.1.1                    | DS1404025-E5                   |
|                             | 8691SF/CPU/256 MB                                 | 8691SF/256 with 256 MB SDRAM installed                                                             | 3.1.1                    | DS1404090-E5                   |
|                             | 8692SF/CPU                                        | Switching fabric                                                                                   | 3.5.6, 3.7.3, 4.0.0      | DS1404065-E5                   |
|                             | 8692SFw/SuperMezz                                 | 8692SF Switch Fabric/CPU with factory-installed Enterprise Enhanced CPU Daughter Card (SuperMezz). | 4.1.0                    | DS1404066-E5                   |
|                             | Enterprise Enhanced CPU Daughter Card (SuperMezz) | Optional daughter card for the 8692 SF/CPU                                                         | 4.1.0                    | DS1411025-E5                   |
| <b>Power supplies</b>       |                                                   |                                                                                                    |                          |                                |
|                             | 8004AC                                            | 850 W AC                                                                                           | 3.1.2                    | DS1405x08                      |

**Table 7**  
**Chassis, power supply, and SF/CPU compatibility (cont'd.)**

| Item                |                           |                                                                                                                                                  | Minimum software version | Part number  |
|---------------------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|--------------|
| <b>Chassis</b>      |                           |                                                                                                                                                  |                          |              |
|                     | 8004DC                    | 850 W DC                                                                                                                                         | 3.1.2                    | DS1405007    |
|                     | 8005AC                    | 1462 W AC                                                                                                                                        | 4.0.0                    | DS1405012    |
|                     | 8005DI AC                 | 1492 W Dual input AC                                                                                                                             | 5.0                      | DS1405016-E6 |
|                     | 8005DC                    | 1462 W DC                                                                                                                                        | 4.0.x                    | DS1405011    |
| <b>Upgrade kits</b> |                           |                                                                                                                                                  |                          |              |
|                     | 256 MB SF/CPU upgrade kit | The 8691 SF/CPU must be upgraded to 256 MB with Software Release 3.5 or later. This memory upgrade is required for the software to run properly. | 3.5.0                    | DS1411016    |

**Table 8**  
**Module and component compatibility**

| Modules and components  |                                                   |                                                                                                                                       | Minimum software version | Part number |
|-------------------------|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------|-------------|
| <b>Security modules</b> |                                                   |                                                                                                                                       |                          |             |
|                         | 8661 SSL Acceleration Module (SAM)                | High Performance SSL Acceleration Module secures web-based applications and business communications                                   | 3.3.1 <sup>2</sup>       | DS1404070   |
|                         | 8660 Service Delivery Module Firewall 1 (SDM FW1) | The 8660 SDM is a combination of dedicated hardware and software that addresses the needs for security, performance, and ease of use. | 3.7.6                    | DS1404104   |
|                         | 8660 Service Delivery Module Firewall 2 (SDM FW2) |                                                                                                                                       | 3.7.6                    | DS1404081   |
|                         | 8660 Service Delivery Module Firewall 4 (SDM FW4) |                                                                                                                                       | 3.7.6                    | DS1404080   |

**Table 8**  
**Module and component compatibility (cont'd.)**

| Modules and components    |                                                                                         |                                             | Minimum software version | Part number |
|---------------------------|-----------------------------------------------------------------------------------------|---------------------------------------------|--------------------------|-------------|
|                           | 8660 Service Delivery Module Threat Protection System (SDM TPS4)                        |                                             | 4.1.0                    | DS1404082   |
|                           | 8660 Combo Service Delivery Module Firewall 2 Threat Protection System 2 (SDM FW2/TPS2) |                                             | 4.1.0                    | DS1404086   |
|                           | 8660 Combo Service Delivery Module Firewall 1 Threat Protection System 1 (SDM FW1/TPS1) |                                             | 4.1.0                    | DS1404087   |
|                           | 8660 SDM spare disk drive                                                               | Replacement part                            | 3.7.6                    | DS1411023   |
|                           | 8660 SDM spare PrPMC                                                                    | Replacement part                            | 3.7.6                    | DS1411024   |
| <b>Layer 4-7 module</b>   |                                                                                         |                                             |                          |             |
|                           | Web Switching Module (WSM)                                                              | 4-port Gigabit Ethernet SX or 10/100BASE-TX | 3.1.3, 3.2.1, 3.3.0      | DS1404045   |
| <b>Ethernet E modules</b> |                                                                                         |                                             |                          |             |
|                           | 8608GBE module                                                                          | 8-port Gigabit Ethernet GBIC                | 3.1.1                    | DS1404038   |
|                           | 8608GTE module                                                                          | 8-port Gigabit Ethernet 1000BASE-TX         | 3.1.1                    | DS1404044   |
|                           | 8608SXE module                                                                          | 8-port Gigabit Ethernet 1000BASE-SX         | 3.1.1                    | DS1404036   |
|                           | 8616SXE module                                                                          | 16-port Gigabit Ethernet 1000BASE-SX        | 3.1.0                    | DS1404011   |
|                           | 8616GTE module                                                                          | 16-port Gigabit Ethernet 1000BASE-TX        | 3.3.0                    | DS1404034   |
|                           | 8624FXE module                                                                          | 24-port 100BASE-FX                          | 3.1.1                    | DS1404037   |
|                           | 8648TXE module                                                                          | 48-port 10/100BASE-TX                       | 3.1.1                    | DS1404035   |
|                           | 8632TXE module                                                                          | 32-port 10/100BASE-TX (2 GBICs)             | 3.1.2                    | DS1404024   |

**Table 8**  
**Module and component compatibility (cont'd.)**

| Modules and components                    |                  |                                                                              | Minimum software version | Part number  |
|-------------------------------------------|------------------|------------------------------------------------------------------------------|--------------------------|--------------|
| <b>Ethernet M modules</b>                 |                  |                                                                              |                          |              |
|                                           | 8608GBM module   | 8-port Gigabit Ethernet GBIC                                                 | 3.3.0                    | DS1404059    |
|                                           | 8608GTM module   | 8-port Gigabit Ethernet 1000TX                                               | 3.3.0                    | DS1404061    |
|                                           | 8632TXM module   | 32-port 10/100BASE-TX (2 GBICs)                                              | 3.3.0                    | DS1404055    |
|                                           | 8648TXM module   | 48-port 10/100BASE-TX                                                        | 3.3.0                    | DS1404056    |
| <b>Ethernet R modules</b>                 |                  |                                                                              |                          |              |
|                                           | 8630GBR module   | 30-port Gigabit Ethernet SFP                                                 | 4.0.0                    | DS1404063    |
|                                           | 8648GTR module   | 48-port 10/100/1000BASE-TX                                                   | 4.0.x                    | DS1404092    |
|                                           | 8683XLR module   | 3-port XFP (10.3125 Gb/s LAN PHY)                                            | 4.0.0                    | DS1404101    |
|                                           | 8683XZR module   | 3-port XFP (10.3125 Gb/s LAN PHY and 9.953 Gb/s WAN PHY)                     | 4.1.0                    | DS1404064    |
| <b>Ethernet RS modules</b>                |                  |                                                                              |                          |              |
|                                           | 8648GTRS         | 48-port 10/100/1000 Mbps copper ports                                        | 5.0.0                    | DS1404110-E6 |
|                                           | 8612XLRS         | 12-port 10 GbE LAN module                                                    | 5.0.0                    | DS1404097-E6 |
|                                           | 8634XGRS         | 24 100/1000 Mbps SFP ports<br>2 XFP ports<br>8 10/100/1000 Mbps copper ports | 5.0.0                    | DS1404109-E6 |
|                                           | 8648GBRS         | 48 100/1000 Mbps SFP ports                                                   | 5.0.0                    | DS1404102-E6 |
| <b>Asynchronous Transfer Mode modules</b> |                  |                                                                              |                          |              |
|                                           | 8672ATME module  | ATME module                                                                  | 3.1.1                    | DS1304008    |
|                                           | 8672ATMM module  | ATMM module                                                                  | 3.3.0                    | DS1304009    |
| <b>ATM module components</b>              |                  |                                                                              |                          |              |
|                                           | DS-3 MDA         | 2-port 75 ohm coaxial                                                        | 3.3.0                    | DS1304002    |
|                                           | OC-12c/STM-4 MDA | 1-port multimode fiber (MMF)                                                 | 3.1.0, 3.1.1, 3.3.0'     | DS1304004    |
|                                           | OC-12c/STM-4 MDA | 1-port single mode fiber (SMF)                                               | 3.1.0, 3.1.1, 3.3.0'     | DS1304005    |
|                                           | OC-3c/STM-1 MDA  | 4-port MMF                                                                   | 3.1.0, 3.1.1, 3.3.0'     | DS1304006    |

**Table 8**  
**Module and component compatibility (cont'd.)**

| Modules and components                            |                           |                                                 | Minimum software version | Part number                  |
|---------------------------------------------------|---------------------------|-------------------------------------------------|--------------------------|------------------------------|
|                                                   | OC-3c/STM-1 MDA           | 4-port SMF                                      | 3.1.0, 3.1.1, 3.3.0'     | DS1304007                    |
| <b>Packet over SONET modules</b>                  |                           |                                                 |                          |                              |
|                                                   | 8683POSM module           | POSM module                                     | 3.3.0                    | DS1404060                    |
| <b>Packet over SONET Media Dependent Adapters</b> |                           |                                                 |                          |                              |
|                                                   | OC-3c/STM-1 MDA           | 2-port MMF                                      | 3.1.0, 3.1.1, 3.3        | DS1333003                    |
|                                                   | OC-3c/STM-1 MDA           | 2-port SMF                                      | 3.1.0, 3.1.1, 3.3        | DS1333004                    |
|                                                   | OC-12c/STM-4 MDA          | 1-port MMF                                      | 3.1.0, 3.1.1, 3.3        | DS1333001                    |
|                                                   | OC-12c/STM-4 MDA          | 1-port SMF                                      | 3.1.0, 3.1.1, 3.3        | DS1333002                    |
| <b>Gigabit Interface Converters</b>               |                           |                                                 |                          |                              |
|                                                   | 1000BASE-SX GBIC          | 850 nanometers (nm) SC duplex                   | 3.0.0                    | AA1419001-E5                 |
|                                                   | 1000BASE-LX GBIC          | 1300 nm SC duplex                               | 3.0.0                    | AA1419002-E5                 |
|                                                   | 1000BASE-XD GBIC          | SC duplex SMF                                   | 3.0.0                    | AA1419003-E5                 |
|                                                   | 1000BASE-ZX GBIC          | 1550 nm, SC duplex SMF                          | 3.0.0                    | AA1419004-E5                 |
|                                                   | 1000BASE-EX CWDM APD GBIC | SC connector, 1470 to 1610 nm                   | 3.1.4                    | AA1419017-E5 to AA1419024-E5 |
|                                                   | 1000BASE-T GBIC           | Category 5 copper unshielded twisted pair (UTP) | 3.5.0                    | AA1419041-E5                 |
| <b>Small form factor pluggable transceivers</b>   |                           |                                                 |                          |                              |
|                                                   | 1000BASE-SX SFP           | 850 nm LC connector                             | 4.0.0                    | AA1419013-E5                 |
|                                                   | 1000BASE-SX SFP           | 850 nm MT-RJ connector                          | 4.0.0                    | AA1419014-E5                 |
|                                                   | 1000BASE-LX SFP           | 1310 nm LC connector                            | 4.0.0                    | AA1419015-E5                 |
|                                                   | 1000BASE-XD CWDM SFP      | From 1470 nm to 1610 nm LC connector            | 4.0                      | AA1419025-E5 to AA1419032-E5 |

**Table 8**  
**Module and component compatibility (cont'd.)**

| Modules and components                                              |                               |                                                                  | Minimum software version | Part number                  |
|---------------------------------------------------------------------|-------------------------------|------------------------------------------------------------------|--------------------------|------------------------------|
|                                                                     | 1000BASE-ZX CWDM SFP          | From 1470 nm to 1610 nm LC connector                             | 4.0                      | AA1419033-E5 to AA1419040-E5 |
|                                                                     | 1000BASE-T SFP                | Category 5 copper unshielded twisted pair (UTP), RJ-45 connector | 4.0.0                    | AA1419043-E6                 |
|                                                                     | 1000BASE-SX SFP               | 850 nm DDI LC connector                                          | 5.0                      | AA1419048-E6                 |
|                                                                     | 1000BASE-LX SFP               | 1310 nm DDI LC connector                                         | 5.0                      | AA1419049-E6                 |
|                                                                     | 1000BASE-XD SFP               | 1310 nm DDI LC connector                                         | 5.0                      | AA1419050-E6                 |
|                                                                     | 1000BASE-XD SFP               | 1550 nm DDI LC connector                                         | 5.0                      | AA1419051-E6                 |
|                                                                     | 1000BASE-ZX SFP               | 1550 nm DDI LC connector                                         | 5.0                      | AA1419052-E6                 |
|                                                                     | 1000BASE-XD CWDM SFP          | From 1470 nm to 1610 nm, DDI                                     | 5.0                      | AA1419053-E6 to AA1419060-E6 |
|                                                                     | 1000BASE-ZX CWDM SFP          | From 1470 nm to 1610 nm, DDI                                     | 5.0                      | AA1419061-E6 to AA1419068-E6 |
|                                                                     | 1000BASE-BX bidirectional SFP | 1310 nm, single fiber LC                                         | 4.1.0                    | AA1419069-E6                 |
|                                                                     | 1000BASE-BX bidirectional SFP | 1490 nm, single fiber LC                                         | 4.1.0                    | AA1419070-E6                 |
|                                                                     | 1000BASE-EX                   | 1550 nm, up to 120 km                                            | 5.0                      | AA1419071-E6                 |
| <b>10 Gigabit Ethernet Small form factor pluggable transceivers</b> |                               |                                                                  |                          |                              |
|                                                                     | 10GBASE-LR/LW XFP             | 1-port 1310 nm SMF, LC connector                                 | 4.0.0                    | AA1403001-E5                 |
|                                                                     | 10GBASE-ER/EW XFP             | 1-port 1550 nm SMF, LC connector                                 | 4.0.x                    | AA1403003-E5                 |
|                                                                     | 10GBASE-SR/SW XFP             | 1-port 850 nm MMF, LC connector                                  | 4.0.0                    | AA1403005-E5                 |

**Table 8**  
**Module and component compatibility (cont'd.)**

| Modules and components |                   |                           | Minimum software version | Part number  |
|------------------------|-------------------|---------------------------|--------------------------|--------------|
|                        | 10GBASE-ZR/ZW XFP | 1550 nm SMF LC connector  | 4.1.0                    | AA1403006-E5 |
|                        | 10GBASE-LRM XFP   | Up to 220 m over MMF, DDI | 5.0.0                    | AA1403007-E6 |

## Module and mode compatibility

R mode allows the operation of R and RS module-specific features. The following table shows which modules (E, M, and R modules) are enabled, depending on the operation mode that you configure and the configuration of your system.

**Table 9**  
**Module and mode interoperability**

|                           |                | Module types            |   |   |
|---------------------------|----------------|-------------------------|---|---|
| Chassis configuration     | Operation mode | R and RS                | M | E |
|                           |                | (e=enabled, d=disabled) |   |   |
| Same-type module chassis  | Default        |                         |   | e |
|                           | M              |                         | e |   |
|                           | R              | e                       |   |   |
| Mixed-type module chassis | Default        | e                       | e | e |
|                           | M              | e                       | e | d |
|                           | R              | e                       | d | d |

## Module type and feature compatibility

The following table shows the supported features for each module.

| Feature                   | Module types |   |   | Comment                                                             |
|---------------------------|--------------|---|---|---------------------------------------------------------------------|
|                           | R and RS     | M | E |                                                                     |
| ACL IN L2 to L4           | yes          |   |   | Ingress filtering ACT, ACL, ACE based                               |
| ACL OUT L2 to L4          | yes          |   |   | Egress filtering ACT, ACL, ACE based                                |
| ACL pattern matching      | yes          |   |   | Pattern match filtering for ingress and egress                      |
| Ingress policing L2 to L4 | yes          |   |   | 450 policers per LANE (10x1 Gb/s, 1x10 Gb/s), total 10 800 policers |

| Feature                 | Module types |     |     | Comment                                                                                                                             |
|-------------------------|--------------|-----|-----|-------------------------------------------------------------------------------------------------------------------------------------|
|                         | R and RS     | M   | E   |                                                                                                                                     |
| Egress shaping L2 to L4 | yes          |     |     | per port shapers, per queue shapers                                                                                                 |
| Egress queues 64        | yes          |     |     | 640 queues per LANE (10x1 Gb/s, 1x10 Gb/s), total 15 360 queues                                                                     |
| SMLT over 10 Gbps       | yes          | yes | yes | in any mode available (8683XLR required)                                                                                            |
| Port mirroring          | yes          | yes | yes | Ingress R and RS modules: up to 20 ports per LANE<br>Egress R module: 1 port per LANE;<br>Egress RS module, up to 20 ports per LANE |
|                         |              |     |     |                                                                                                                                     |
| SVLAN                   | no           | yes | yes |                                                                                                                                     |
| IPX routing             | no           | yes | yes |                                                                                                                                     |
| Classic filter L3-L4    | no           | yes | yes |                                                                                                                                     |
| Classic policing L3-L4  | no           | yes | yes |                                                                                                                                     |

## Module and chassis compatibility and performance considerations

All existing Ethernet Routing Switch 8600 hardware is supported in Release 4.1 and later with the exception of 8100 modules and systems. 8100 modules and systems are not supported with any 4.x or later release. The 8003 chassis is supported with any 4.x or 5.x Release, but the 8003 chassis does not support the 8692 SF/CPU, and therefore does not support R or RS modules.

The new RS modules are compatible with existing 8010 and 8006 chassis.

In older chassis (those shipped before 2005), there is a difference between Standard and High Performance slots. In these chassis, an R or RS module installed in a Standard slot delivers increased port density. An R or RS module installed in a High Performance slot delivers increased port density and increased performance. Chassis manufactured in 2005 and later do not have this limitation, and have full high-performance slot support.

In older chassis, R and RS modules inserted in slots 2 to 4 and slots 7 to 9 of the 8010 10-slot chassis, and slots 2 to 4 of the 8006 6-slot chassis, always operate at high performance. R modules inserted into slot 1 and slot 10 of the 8010 chassis, and slot 1 of the 8006 chassis, can operate at high performance, but operate at standard performance depending on chassis revision (for more information about identifying chassis, see the

following section). For information about relative performance per slot with two fabrics installed in existing 8010, 8010co, and 8006 chassis, see the following table.

**Table 10**  
**Pre-2005 8010, 8010co, and 8006 chassis performance**

| Module   | Standard slot (Slots 1 and 10) full duplex | High Performance slot (Slots 2 to 4, Slots 7 to 9) full duplex |
|----------|--------------------------------------------|----------------------------------------------------------------|
| E and M  | 16 Gbps                                    | 16 Gbps                                                        |
| 8630GBR  | 16 Gbps                                    | 60 Gbps                                                        |
| 8683XLR  | 16 Gbps                                    | 60 Gbps                                                        |
| 8648GTR  | 16 Gbps                                    | 32 Gbps                                                        |
| 8683XZR  | 16 Gbps                                    | 60 Gbps                                                        |
| 8612XLRS | 16 Gbps                                    | 60 Gbps                                                        |
| 8648GTRS | 16 Gbps                                    | 40 Gbps                                                        |
| 8648GBRS | 16 Gbps                                    | 60 Gbps                                                        |
| 8634XGRS | 16 Gbps                                    | 60 Gbps                                                        |

If you place an R or RS module into a Standard slot of a non-high performance chassis, you receive the following message:

For maximum performance, Nortel recommends placing R modules in Slots 2 to 4 or 7 to 9 as available. Please refer to release notes for additional details.

### High Performance chassis

A chassis revision with an upgraded High Performance Backplane is available. The High Performance chassis is compatible with existing E, M, R, and RS modules.

Identify the High Performance Backplane by using the CLI or NNCLI. Use the CLI command **show sys info** or the NNCLI command **show sys-info** to show the chassis revision number. The HwRev field indicates if the chassis is High Performance or Standard. The following table provides the Hardware Revision details for each chassis model. For more information, see the Technical Tip *Identifying the new Ethernet Routing Switch 8600 Chassis* (TT-0507501A) on the Nortel Technical Support Web site.

**Table 11**  
**Chassis hardware revision**

| Chassis model | Hardware Revision                                | H/W Config    |
|---------------|--------------------------------------------------|---------------|
| 8006          | 05 or greater indicates high performance chassis | 02 or greater |
| 8010          | 06 or greater indicates high performance chassis | 02 or greater |
| 8010co        | 05 or greater indicates high performance chassis | 02 or greater |

Customers requiring High Performance Mode for all slots on an older Ethernet Routing Switch 8600 chassis can have their existing chassis exchanged and reworked. Call 1-800-4NORTEL and order service part number N0060024. The list price for this chassis re-work is US \$2000.00 for each chassis, and an advanced replacement unit is provided.

## Switch clustering topologies and interoperability with other products

When the Ethernet Routing Switch 8600 is used with other Ethernet Routing Switch products, the switch clustering bridging, unicast routing, and multicast routing configurations vary with switch type. Nortel recommends that you use the supported topologies and features when you perform inter-product switch clustering. For more information, see *Switch Clustering (SMLT/SLT/RSMLT/MSMLT) Supported Topologies and Interoperability with ERS 8600 / 5500 / 8300 / 1600* (NN48500-555), available on the Nortel Technical Support Web site.

For specific design and configuration parameters, see *Converged Campus Technical Solutions Guide* (NN48500-516) and *Switch Clustering using Split-Multilink Trunking (SMLT) Technical Configuration Guide* (NN48500-518).

## SF/CPU protection and loop prevention compatibility

Nortel recommends several best-practice methods for loop prevention, especially in any Ethernet Routing Switch 8600 Switch cluster environment. For more information about loop detection and compatibility for each software release, see *Converged Campus Technical Solution Guide — Enterprise Solution Engineering* (NN48500-516).

## Switch behavior during boot cycle and redundant configuration files

Nortel recommends that you take special care when providing the boot option for your production systems. The Ethernet Routing Switch 8600 provides three boot configuration file choices, as well as a backup configuration file choice for each configuration file choice.

The default boot sequence directs the switch to look for its image and configuration files first on the PCMCIA card, then in the onboard flash memory, and then from a server on the network. The switch first checks for /pcmcia/pcmbboot.cfg and then checks for /flash/boot.cfg.

The PCMCIA card is the primary source for the files; the onboard flash memory is the secondary source; and the network server is the tertiary source. These source and file name definitions are in the boot configuration file. The boot source order is configurable.

The config.cfg file stores the configuration of the Ethernet Routing Switch 8600 and its modules. This is the default configuration file. You can specify a different configuration file for the switch to use for the boot process.

For more details about boot sources, see *Nortel Ethernet Routing Switch 8600 Administration* (NN46205-605).

In normal operation, Nortel recommends that the primary configuration file is saved on the /flash drive, and that the primary backup configuration file is saved on the /pcmcia drive. Using this configuration, if one file or drive gets corrupted, the switch can still boot from the other file or drive. When you change configuration files, Nortel further recommends that you save the last known good configuration using the secondary choice option.



### CAUTION

#### Risk of network outage

If a switch cannot access a valid configuration file, it will fall into default configuration mode, which can cause a network outage.

Ensure that a valid configuration and a backup configuration file are always available.

### ATTENTION

If you want to store only one simple backup configuration file, Nortel recommends that you use a default backup configuration file with the following information (only) included:

```
config ethernet 1/1-10/48 state disable
```

This ensures that all ports remain disabled if the backup configuration file is loaded for any reason.

The information in the following table describes how the switch behaves in different boot situations. If a configuration file is unspecified, this means that the **config bootconfig choice** command was not provided for the file. The switch action column describes the expected behavior in both CLI and NNCLI modes, unless otherwise specified.

**Table 12**  
**Switch behavior during boot cycle**

| Parameters                                                                                                                                                                            | Switch action                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A configuration file is not specified.<br>The config.cfg file is present on the flash drive.                                                                                          | The switch boots config.cfg                                                                                                                                                                                                                                                                                                                                                |
| The primary configuration file is specified.<br>The configuration file is present on the flash drive.                                                                                 | The switch boots the specified configuration file.                                                                                                                                                                                                                                                                                                                         |
| The primary configuration file is specified.<br>The configuration file is not present on the flash drive.                                                                             | The switch boots with factory defaults (if <b>config boot flags verify-config</b> is <b>true</b> , and a backup configuration is not specified).                                                                                                                                                                                                                           |
| The primary configuration file is specified.<br>The configuration file on the flash drive has a bad command.                                                                          | The switch boots with factory defaults (if <b>config boot flags verify-config</b> is <b>true</b> , and a backup configuration is not specified).                                                                                                                                                                                                                           |
| The primary configuration file is specified.<br>The configuration file on the flash drive has a bad command.<br>The backup configuration file is specified, but it has a bad command. | The switch fails the first configuration file, and boots the second configuration file, ignoring the bad command.                                                                                                                                                                                                                                                          |
| The switch is configured to boot with factory defaults.                                                                                                                               | The switch boots with factory defaults.                                                                                                                                                                                                                                                                                                                                    |
| The boot.cfg file is corrupt.                                                                                                                                                         | In CLI mode: The switch fails to load the boot.cfg file and creates a new boot.cfg file with a default boot configuration.<br>In NNCLI mode: The switch fails to load the boot.cfg file and creates a new boot.cfg file with a default boot configuration. The switch comes up in CLI mode, which is the correct behavior because the NNCLI mode flag is false by default. |

### Configuring primary, secondary, and tertiary boot sources

Configure the boot sources so that the switch uses proper files from which to boot.

| Step    | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1       | <p>To change the runtime configuration file locations, use the following command:</p> <pre>config bootconfig choice &lt;primary   secondary   tertiary&gt; [config-file &lt;file&gt;   backup-config-file &lt;file&gt;   image-file &lt;file&gt;]</pre> <p>For example, to specify the configuration file in flash memory as the primary, use the following command:</p> <pre>ERS-8610:6# config bootconfig choice primary config-file /flash/config.cfg</pre>                                                                                                                                                                                                                     |
| 2       | <p>To set the location for the I/O module driver image for the BootStrap protocol:</p> <pre>config bootconfig bootp image-name &lt;image-name&gt; &lt;slot-number&gt;</pre> <pre>config bootconfig bootp secondary-image-name &lt;image-name&gt; &lt;slot-number&gt;</pre> <p>For example, to specify an R module driver for slot 2 in flash memory, use the following command:</p> <pre>ERS-8610:6#config bootconfig bootp /flash/p80j50 xx.dld 2</pre> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>ATTENTION</b><br/>Nortel recommends that you store .dld files in flash memory, and that you always set the image-name to default.</p> </div> |
| 3       | <p>To set the boot source location for the SuperMezz image:</p> <pre>config bootconfig mezz-image image-name &lt;image-name&gt;</pre> <p>For example:</p> <pre>ERS-8610:6#config bootconfig mezz-image image-name /flash/p80m50xx.img</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| --End-- |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

The following example configures the primary and secondary sources as per Nortel recommendations.

| Step | Action                                                                                                                                    |
|------|-------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Configure the primary configuration file choices:</p> <pre>config bootconfig choice primary config-file /flash/primaryconfig.cfg</pre> |

```
config bootconfig choice primary backup-config-file
/pcmcia/primaryconfig.cfg

2 Configure the secondary configuration file choices:

config bootconfig choice secondary config-file
/flash/secondaryconfig.cfg

config bootconfig choice secondary backup-config-file
/pcmcia/secondaryconfig.cfg
```

---

--End--

---

## MPLS considerations

The MPLS maximum transmission unit (MTU) is dynamically provisioned (1522 or 1950 bytes) and it supports jumbo frames (9000 bytes). Packets that exceed the MTU are dropped. The allowed data CE frame size is MTU size minus MPLS encapsulation (header) size. For control frames (for example, LDP) the frame size is 1522 or 1950 bytes.

For the Ethernet Routing Switch 8600, the MPLS RSVP LSP Retry Limit is infinite by design (a setting of zero means infinite). When the limit is infinite, should a Label Switched Path (LSP) go down, it is retried using exponential backoff. The Retry Limit is not configurable.

In scaled environments, if MPLS LDP sessions flap and CPU utilization increases, then the default Hello Hold Timer of 60 seconds may not be long enough. If this situation occurs, Nortel recommends that you increase the Hold Timer to 120 or 180 seconds.

## SNMP considerations

Release 5.0 introduces changes to SNMP behavior. SNMP is configured differently in the NNCLI than in the CLI. Auto-generation of several parameters and command structure changes means that several configuration procedures are no longer required in the NNCLI. For more information:

- Refer to CR Q01858455.
- For SNMP trap changes, see the NNCLI SNMP trap configuration section in *Nortel Ethernet Routing Switch 8600 Troubleshooting* (NN46205-703).
- For SNMP community-based changes, see *Nortel Ethernet Routing Switch 8600 Administration* (NN46205-605) in this document.

## VLACP considerations

Nortel recommends the following:

- The best practice standard settings for VLACP are a short timer of no less than 500 milliseconds and a time-out scale of 5. Both faster timers and lower time-out scales are supported, but if any VLACP flapping occurs, increase the short timer and the time-out scale to their recommended values: 5 and 500, respectively.
- Do not use VLACP on any configured LACP MLTs because LACP provides the same functionality as VLACP for link failure. VLACP and LACP running on the same link is not supported.
- VLACP with short timers is not supported on interswitch trunk (IST) links. Use only long timers. For IST links, Nortel recommends that you do not set the VLACP long periodic timer to less than 30 seconds.
- Although the software configuration supports VLACP short timers of less than 30 ms, using values less than 30 ms is not supported in practice. The shortest (fastest) supported VLACP timer is 30 ms with a timeout of 3, which is used to achieve sub-100 ms failover (see “VLACP sub-100 ms convergence using SuperMezz” (page 78) ). 30 ms timers are not supported in High Availability (HA) mode, and may not be stable in scaled networks. At this time, 30 ms timers are only supported between two Ethernet Routing Switch 8600 switches using SuperMezz. VLACP timer values below 500 ms may also require the SuperMezz option, particularly in scaled environments or when very short (fast) timer values are used.

For more information on VLACP best practices, see *Nortel Ethernet Routing Switch 8600 Planning and Engineering — Network Design* (NN46205-200).

## DVMRP considerations

For Distance Vector Multicast Routing Protocol (DVMRP) configurations of more than 1000 streams, you may have to increase protocol timeouts (for example, OSPF dead interval, and so on). Otherwise, traffic loss can occur.

## SMLT considerations

Software Release 5.1 does not support PIM Multicast Border Router (MBR) functionality over SMLT.

Nortel does not support an additional redundant MLT group between two IST peers.

In a large-scale SMLT network (that is, one with more than 10 000 MAC addresses), Nortel strongly recommends that you operate the system with M, or preferably, R and RS modules. Further, if you use R or RS modules, use R mode. This ensures that the system has enough memory and record space available for proper switch operation.

To improve SMLT failover and recovery behavior for large-scale networks, Nortel has optimized the IST protocol and rearchitected the SMLT state machines. This functionality improvement is mainly targeted for large-scale SMLT networks, where Layer 2 bridging, Layer 3 routing with VRRP BackupMaster or RSMLT, and Multicast SMLT (IGMP and PIM) are required. Using these optimizations, failover and recovery improvements occur for networks in which IST switches handle more than 5000 MAC address table entries, 5000 ARP table entries, and 100 multicast streams.

For proper network behavior, Nortel recommends that you operate both IST switches with either the new or old SMLT architecture. That is, both switches on either end of an IST must use the same SMLT architecture and therefore the same software release version. Interconnected IST pairs may run different software release versions. Release 5.1.0 does not include the SMLT re-architecture.

### RSMLT considerations

In an RSMLT configuration, to ensure peer forwarding when the peer is down, enter save config after the peer information is first learned by both peers, or at any later time when the peer RSMLT information changes.

Whenever the peer RSMLT information changes (for example, from adding or deleting VLANs, changing VLAN IDs, or changing VLAN IP addresses), messages appear in the log indicating a discrepancy between stored information and what the switch is receiving from the peer. For example:

```
CPU6 [03/07/09 01:25:19] IP WARNING Recvd Peer address not same as
stored address for Vlan 544.  Save config for Edge-Support to use
this info on next reboot
CPU6 [03/07/09 01:25:19] IP WARNING Recvd Peer address not same as
stored address for Vlan 536.  Save config for Edge-Support to use
this info on next reboot
CPU6 [03/07/09 01:25:19] IP WARNING Recvd Peer address not same as
stored address for Vlan 535.  Save config for Edge-Support to use
this info on next reboot
```

When the preceding messages appear in the log, if the peer goes down, the switch does not forward the traffic for its peer for the indicated VLANs unless the configuration is saved.

## MSDP considerations

To achieve MSDP SA scalability with maximum MSDP peers, Nortel recommends that you install the 8692 SF/CPU module, with or without SuperMezz.

On a switch equipped with the 8691 SF/CPU and E modules that is scaled to 6144 SAs, MSDP may flap, but only rarely. The flapping can be seen easily when using Jumbo frames. (Q01992856)

## High Availability mode considerations

High Availability mode (also known as HA-CPU) permits the synchronization of configuration and protocol states between the Master and Secondary CPUs.

For Release 5.1, HA-CPU supports the following in Hot Standby mode:

- platform configuration
- Layer 2 protocols: IGMP, STP, MLT, SMLT, ARP, LACP, VLACP
- Layer 3 protocols: RIP, OSPF, VRRP, RSMLT, VRF Lite

HA-CPU supports the following in Warm Standby mode. After failover, these protocols are restarted:

- DVMRP, PIM-SM, PIM-SSM
- BGP
- MPLS
- BFD

Hot Standby mode performs hitless failover, while Warm Standby mode restarts protocols after failover.

A reboot is necessary to make HA-CPU mode active.

HA-CPU does not currently support the following protocols or modules:

- PGM
- IPX RIP/SAP
- Web Switching Module (WSM)
- 8683 POS Module
- 8660 SDM Module
- 8672 ATM Module

In HA mode, Nortel recommends that you do not configure the OSPF Hello timers for less than one second, and the dead router interval for less than 15 seconds (Q01445505).

When using a SuperMezz daughtercard and Classic modules in HA environments, HA failover can take some time. This may cause other side effects, for example, ISTs may go up and down, or traffic loss may occur.

When using an 8691SF/CPU with Classic modules and performing HA-CPU failovers, OSPF neighborhood may be re-established in some situations.

When using the CLI or NNCLI on a dual-CPU system with HA mode enabled, do not enter configuration commands on the Standby SF/CPU. Execute all configuration commands on the Master SF/CPU only.

If you use Layer 3 multicast routing protocols in HA mode, it takes approximately 30 seconds for the routing protocols to re-converge after an HA failover. This behavior occurs because multicast protocols are only partially supported by Release 5.1, and thus PIM and DVMRP are restarted upon failover. When PIM is enabled, it takes the same amount of time (approximately 30 seconds) to come up.

### **HA mode and short timers**

Prior to Release 5.0, protocols that used short timers could bounce (restart) during HA failover. These protocols include VLACP, LACP, VRRP, OSPF, and STP. Release 5.0 and later support fast failover for configurations that use short timers. In Release 5.0 and later:

- all HA configurations that use R or RS modules in R mode, along with the 8692 SF/CPU with SuperMezz, supports protocols with short timers. Fast failover under these conditions is supported.
- all HA configurations that use E or M modules, along with the 8692 SF/CPU with SuperMezz, may restart the protocols upon failover.
- all HA configurations that use E, M, R or RS modules in default, M, or R mode, along with the 8692 SF/CPU without SuperMezz, supports protocols with short timers. Fast failover under these conditions is supported.
- any HA configuration that uses the 8691 SF/CPU may restart the protocols upon failover.

### **PCAP considerations**

Use of the Packet Capture (PCAP) tool is not supported when a SuperMezz daughtercard is used.

## WSM considerations

Release 5.1 supports Web Switching Module WebOS software version 10.0.34.0. The Web Switching Module WebOS software implementation for STP over MLT is fully compatible with Ethernet Routing Switch 8600 Software Release 5.1.

If a Web Switching Module (WSM) is installed, any switches set to either R or M mode do not initialize the WSM.

If you trunk the Web Switching Module with Alteon Stackable 180 Series switches, Spanning Tree Protocol (STP) must be disabled on either the Web Switching Module or the Alteon Stackable to avoid incorrect STP operation.

## Enterprise Switch Manager (ESM) considerations

ESM code release 6.3 provides support for Release 5.1.0 functionality.

## 60 day trial license

You are provided a 60 day trial period of Ethernet Routing Switch 8600, during which you have access to all features. In the trial period you can configure all features without restriction, including system console and log messages. The switch logs trial period messages even if no license features are used or tested during the trial period. If any valid license is loaded on the switch at any time, none of the trial period messages are logged. At the end of the trial period, a message appears notifying the user that the trial period has expired.

For additional information about trial licenses, see *Nortel Ethernet Routing Switch 8600 Administration* ((NN46205-605)).

## Advanced filter guidelines

Use the following guidelines when you configure advanced Layer 2 to Layer 7 filters for R or RS module ports or for VLANs with R or RS module ports in them.

- Always use an ACT with only the proper attributes selected. If you must add ACEs with attributes that are not in the original ACT, you must create a new ACL associated with the new ACT.
- For filter optimization reasons, when you have multiple ACEs that perform the same task (for example: deny or allow IP addresses, or UDP/TCP-based ports), you can configure one ACE to perform the task with either multiple address entries, or address ranges, or a combination of both. You can use this one ACE instead of using multiple ACEs.

For R and RS module ACLs, a maximum of 500 ACEs are supported. This maximum may not be achievable depending on the type of attributes used within an ACE. Since there are millions of combinations, note that certain combinations can overextend the system. In these cases, to help ensure stable system operation, reduce the number of ACEs and follow the previous guidelines.

**CAUTION****Risk of module reset or improper load of configuration file**

If the following messages appear on the console or in the log file, it is likely that there is a specific problematic combination of ACEs configured within an ACL. Such combinations are very unlikely to occur, but if you see these messages, first reduce the number of ACEs within the ACL until the messages stop. Next, contact Nortel Technical Support. Support will attempt to find a combination that does not cause this situation, and will provide the required filtering capabilities.

```
CPU5 [05/23/06 10:51:08] COP-SW ERROR Slot 3:
ercdAddCollapseBin: rcdRspMalloc failed for
INGRESS RSP memory allocation
```

```
CPU5 [05/23/06 10:51:08] COP-SW ERROR Slot 3:
ercdGetCollapseNode: collapse node creation
failed.
```

```
CPU5 [05/23/06 10:51:08] COP-SW ERROR Slot 3: erc
dFilterRdxResultUpdate: ercdGetCollapseNode()
Failed !!
```

## Supported standards, RFCs, and MIBs

For information about supported standards, RFCs, and MIBs, see the Appendices in *Nortel Ethernet Routing Switch 8600 Planning and Engineering — Network Design* (NN46205-200).

## Supported traps and notifications

For a complete list of log messages generated by Ethernet Routing Switch 8600 Software Release 5.1, see *Nortel Ethernet Routing Switch 8600 Logs Reference* (NN46205-701).

For a complete list of SNMP traps generated by Ethernet Routing Switch 8600 Software Release 5.1, see *Nortel Ethernet Routing Switch 8600 Troubleshooting* (NN46205-703).

## Device Manager installation requirements

Device Manager is supported on Windows, Solaris, and Linux.

## Windows

The minimum system requirements for installing Device Manager on Microsoft Windows Vista, Windows 2000, Windows 2003 (Professional) and Windows XP are:

- 512 MB of RAM
- 400 MB space on hard drive

## Solaris

Solaris™/Sun™OS 2.8, 2.9, and 2.10/5.8, 5.9, and 5.10 Device Manager requires Solaris 8 as a minimum requirement. The minimum system requirements for installing Device Manager on Solaris are:

- 512 MB RAM
- 400 MB space on hard drive

## Linux

The minimum system requirements for installing Device Manager on Linux are:

- Kernel version 2.2 and above
- 512 MB RAM
- 400 MB space on hard drive

## Software Release 3.7 support

Software changes and bug fixes are no longer being made to Release 3.7. For more information, see *TECHNICAL SUPPORT PRODUCT LIFECYCLE BULLETIN — Software Support for ERS8600* (2008008618, Rev 1).



## Resolved issues

This section details all issues resolved for Release 5.1.

### Layer 2 switching resolved issues

**Table 13**  
**Layer 2 resolved issues**

| CR references | Description                                                                                                                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01731929-01  | When there are more than eight LACP link members associated with an MLT and the active MLT ports are disabled, the MLT port membership information is not updated, and the standby links do not take over.<br><br>Workaround: Disable and enable the MLT. |
| Q01667281-02  | If there is an additional MLT in parallel to an IST_MLT between a pair of Dual Core IST switches, the MAC hardware records will now get removed once the fdb entry ages out.                                                                              |
| Q01885431-01  | ERS 8600 is no longer occasionally unstable when multicast mode NLB is configured. This situation is explained in CSB 2008008980.                                                                                                                         |
| Q01944595-04  | MLT ports running RSTP no longer experience traffic forwarding issues when individual MLT ports go up or down.                                                                                                                                            |
| Q01887792     | MSTP will now properly send BPDUs on HA failover and will no longer cause instability in the network.                                                                                                                                                     |
| Q01921984     | MSTP will now converge properly when it contains ports that are not part of the CIST instance.                                                                                                                                                            |
| Q01964667-01  | After the power off/on of an 8600 in an RSTP setup, packets are forwarded to a link and discarded, even though the FDB entry is pointing to the correct interface. The result is communication loss in Layer-2. This problem is resolved.                 |
| Q01867402-01  | For Classic fiber modules, the SMLT status will no longer be incorrectly displayed as "smlt" for MLTs which are down.                                                                                                                                     |
| Q01887792-01  | MSTP will now properly send BPDUs on HA failover and will no longer cause instability in the network.                                                                                                                                                     |
| Q01783263-01  | Ports RSTP/MSTP related configuration will now be retained after a removal and re-insertion of an I/O module.                                                                                                                                             |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01887792-01  | Doing a CPU switchover on an ERS8600 with HA mode and MSTP, when a device is connected with two links, no longer disables MSTP nor causes a loop.                                                                                                                                                                                                                                                          |
| Q01783263-01  | Unplugging an I/O card and plugging it back in no longer changes the status of the port RSTP.                                                                                                                                                                                                                                                                                                              |
| Q01978551-01  | The MSTP: Pathcost no longer changes to the default value when a card-reset is performed.                                                                                                                                                                                                                                                                                                                  |
| Q01422918     | In Split MultiLink Trunk configurations, the clients can lose connectivity after one of the aggregation switches boots up. The issue occurs when 100 or more VLANs are configured on the aggregation switches. The timeout period increases with the number of configured VLANs.<br>Workaround: if using VRRP, set the VRRP hold-down timer to 2 to 3 minutes, or use RSMLT with an RSMLT hold-down timer. |
| Q01669984-01  | The MAC FDB table and system hardware records may not be not completely cleared from a remote IST peer aggregation switch when a single-homed local SMLT link goes down.<br>When the link comes back up the records are relearned.                                                                                                                                                                         |
| Q01791061     | When you use the command <code>show ip rsmlt info</code> , the RSMLT status is shown as up even when RSMLT is disabled on the switch, and the value of the hold-down timer is not properly updated on the RSMLT peers. When the next time hold-down timer counts it starts from the correct value. These display issues do not affect functionality.                                                       |
| Q01863386     | If you create a VLAN with a MAC offset of 500, 501, 502 and 502, an error message is displayed. This is the correct behavior.<br>Workaround: Do not use MAC offsets of 500, 501, 502, or 503.                                                                                                                                                                                                              |
| Q01677505     | When an RSMLT is deleted from an aggregate switch, the far end switch should show the RSMLT as down, but it is shown as up. This is a display issue only; functionality is not affected.                                                                                                                                                                                                                   |

## Switch management resolved issues

**Table 14**  
**Switch management resolved issues**

| CR references | Description                                                                                                                                                                                                                                     |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01983973-01  | 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.<br>These log messages will have format: CPU5 [01/30/09 16:53:19] SNMP INFO SWITCH EXTERNALLY REBOOTED |
| Q01942511     | Enabling of the multicast flooding feature (config ip arp multicast-mac-flooding enable) now syncs to the slave when the ERS8600 is operating in HA mode.                                                                                       |
| Q01913996     | Slow response with JDM/SNMP via in-band management is no longer seen.                                                                                                                                                                           |
| Q02001015-01  | The issue where SSH locks up, blocking further connections to the node, and JDM or console access is lost after attempting to run any output command, has been resolved.                                                                        |

| CR references                        | Description                                                                                                                                         |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01492490,<br>01476296,<br>Q01475131 | FTP, RLOGIN, and TFTP are not supported for IPv6 addresses . However, IPv6 supports HTTP, SSH, TELNET and SNMPv3.                                   |
| Q01879065                            | The switch fails to connect to a RADIUS server if the server was unreachable for a long time, and also all accounting requests to that server fail. |

## Platform resolved issues

Table 15

Platform resolved issues

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01950243-01  | Chassis with a HW Rev field of 50 or above were coming out of manufacturing with the HW Config field set with an invalid value. 8010 and 8010CO chassis displayed invalid log message errors incorrectly indicating Slots 1 and 10 were not High Performance slots, when those slots are High Performance. Loading 5.0.1.0 or higher code into a chassis will reprogram, when needed, the affected field so that the chassis does not report the invalid error log messages.                                                                                                                                                                                     |
| Q01910296-01  | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Q01957190     | On a dual CPU system, the master CPU could fail to communicate with the slave CPU in HA enabled mode. Therefore, when doing a failover to the standby CPU by rebooting the master, the master CPU would then switch to being a slave CPU which would fail to load the configurations correctly. This problem has been resolved.                                                                                                                                                                                                                                                                                                                                  |
| Q01929707     | Grandfathered Advanced Licenses no longer display the MAC address of the chassis, as this is not applicable to licensed usage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Q01923695     | <p>For Classic module connections, after the ARP aging timer expired (default value is 360 minutes), and then the re-boot of a WLAN AP communication of wireless clients through that AP could be affected; this is no longer seen. In association with this event, the following message could be seen:</p> <p><b>SW Warning dpmAddMacAddress: dpmFillMacAddrUpdateMsg failed for mac xxxxxx</b></p> <p>Where mac xxxxxx would be the MAC of the affected wireless client. This SW Warning message has now been change to just be a trace message. This issue was not seen on older code, such as 4.1.6.3, but will no longer been seen if 5.0.1.0 is used.</p> |
| Q01944594-01  | Connection loss in certain RSTP setups can occur after pulling/inserting cables that will change the topology. Though the FDB tables look to be correct, certain packets could be forwarded into a link that is blocked on the remote end. Clearing the FDB table previously resolved the issue until some future reoccurrence. This issue is now resolved.                                                                                                                                                                                                                                                                                                      |

| CR references              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01920749-02               | A manufacturing issue was identified on some 8005AC power supplies in which the power supply was incorrectly programmed. This lead 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. Power supplies should not be RMA requested or the SEEPROM manually reprogrammed. |
| Q01938405-01               | 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.                                                                                                                                                                                                                        |
| Q01910296-01               | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Q01980923                  | Links between an 8600 with 8630GBR and a 4548GT will now be established with no incident.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Q01987771<br>Q02004890     | All known issues with link-flapping on Gig or 10 Gig ports in association with any RS module are now resolved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01987844                  | All known issues with LANE or port lock-ups in association with any RS module, are now resolved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Q01994605,<br>Q01723144-01 | The issue where no more than 503 vlan IP interfaces in r-mode could be created has been resolved. The new limit is 4,000 vlan IP interfaces in chassis built after June 13, 2003 and 1,000 vlan IP interfaces in chassis built before June 13, 2003.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Q01848138-03               | The issue where debug information is not saved when superMezz is enabled, and the software lockup detect feature is triggered, has been resolved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Q01815827-03               | If a Web Switching Module (WSM) is running on a chassis with the boot configuration enhanced-operational-mode flag set to true, the switch may crash if you hot-swap the WSM.<br>Workaround: Do not hot-swap a WSM. Administratively disable the module before you remove it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01822189     | <p>After an SF/CPU crash, the 8692 SF/CPU module can erroneously save .cv files with random names to the PCMCIA card. These files are of size 0 KB. Workaround:</p> <ul style="list-style-type: none"> <li>• Transfer all files from the PCMCIA card to another storage location</li> <li>• Reformat the PCMCIA card using the latest Ethernet Routing Switch 8600 software. Ensure you are running the latest software release, then use the <b>dos-format pcmcia</b> command.</li> <li>• Move all files back to the PCMCIA card</li> </ul> |
| Q01860112     | <p>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.</p>                                                  |
| Q01867508     | <p>When you start SSH for the first time, with "config sys set ssh enable true", the switch generates the DSA keys and stores them. It tries to synchronize the key to the Secondary CPU, but this fails because /flash/.ssh/ does not yet exist on the Secondary. Workaround: Create a .ssh directory in flash on the Master and Secondary.</p>                                                                                                                                                                                             |
| Q01876535     | <p>In the CLI, if you try to configure the boot.cfg file to load the primary image and Mezz image files from the PCMCIA (using the <b>config bootconfig choice primary image-file</b> command), the image file names are not stored in the boot.cfg file, and the previously configured image name is erased. Workaround: Edit the boot.cfg file manually to point to the correct location (PCMCIA).</p>                                                                                                                                     |
| Q01406191     | <p>RADIUS authentication fails when the "Pending Requests" server statistic reaches 10. This is related to the connection being terminated prior to the RADIUS server response.</p>                                                                                                                                                                                                                                                                                                                                                          |
| Q01927767-02  | <p>If a DHCP server is configured wrongly to assign the IP address of the Relay Agent itself to a DHCP client, the DHCP Offer packets sent by the server stay in the 8600 system buffer and are not emptied or de-queued. Buffer Utilization reaches 99% and all connections are dropped.</p>                                                                                                                                                                                                                                                |

## IP Unicast resolved issues

Table 16  
IP Unicast resolved issues

| CR references | Description                                                                                                                                                                                                        |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01908628-01  | <p>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.</p> |

| CR references           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01953415-02            | ERS8600 RSMLT edge support now recovers properly when last active port in RSMLT VLAN is bounced.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Q01862845-01, Q01953743 | ERS8600 in an SMLT environment will no longer become unstable with certain traffic patterns also involving a large number of ARP records.                                                                                                                                                                                                                                                                                                                                                                                       |
| Q01928753               | In earlier RSMLT implementations, the default route was added during the holddown timer in VRF0 with the IST Peer interface as the next-hop. As the default route was not added in the NON-Default VRF, traffic loss was seen. One default route per VRF is now added in all the VRFs which has RSMLT enabled VLANs with the next hop being RSMLT Peer Vlan's IP Interface. Upon expiry of the holddown timer, all default routes added in all the VRFs are removed. These RSMLT added default routes are no longer advertised. |
| Q01885793               | If you have IPv6 OSPF redistribution configurations, and save the configuration file, the redistribution entries may not be saved.<br><br>Workaround: After a save, you may need to manually edit the configuration file and add back the affected IPv6 OSPF redistribution entries.                                                                                                                                                                                                                                            |
| Q01852084               | In Device Manager, you cannot apply new costs for IPv6 static routes. After creation, IPv4 or IPv6 static routes cannot be modified. This behavior is equally applicable for the CLI.<br>Workaround: delete and recreate the static route with the correct cost.                                                                                                                                                                                                                                                                |
| Q01867505               | When the <code>show mpls ldp path</code> command is entered, a console session may stop responding.<br>Workaround: open a telnet session.                                                                                                                                                                                                                                                                                                                                                                                       |
| Q01877407               | When the OSPFv3 redistribution (direct or static route) state is checked multiple times using the CLI ( <code>config ipv6 ospf redistribute direct info</code> ), it appears that the state is changing from enabled to disabled. The state is not changing, this is a display issue only.<br>Workaround: Use Device Manager to check the state.                                                                                                                                                                                |

## MLT/SMLT resolved issues

Table 17  
MLT/SMLT resolved issues

| CR references | Description                                                                                                                                                                                                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01799513-03  | 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. |

| CR references                         | Description                                                                                                                                                                                                                                                             |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01992234-01                          | The issue where invalid ARP entries were being created when a MAC is present in multiple vlans has been resolved.                                                                                                                                                       |
| Q01996928,<br>Q01996937,<br>Q02000507 | The issue where ARP entries were 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, and in which this condition did NOT clear after the IST synch, has been resolved. |

## CLI and NNCLI resolved issues

**Table 18**  
**CLI and NNCLI resolved issues**

| CR references | Description                                                                                                                                                                                                                                                                                                       |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01977384     | ERS 8600 with 8691SF in NNCLI mode saw a double "exit" in the config file. On reboot, the ERS 8600 previously aborted the loading of the config file at the second "exit". This problem is resolved.                                                                                                              |
| Q01838871-01  | 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.                                                |
| Q01595073-01  | Using the NNCLI, the value specified for the VLACP ethertype must be specified in decimal equivalent value. Hexadecimal format is not allowed. For example, an ethertype of "8103" must be entered as "33027".                                                                                                    |
| Q01864112     | In the NNCLI, the stop option does not work in the <b>source &lt;file&gt; stop</b> command.<br>Workaround: source the file without using the stop option.                                                                                                                                                         |
| Q01873397     | When using the NNCLI for ARP configuration in Interface Configuration mode, the loop-detect commands do not operate correctly.                                                                                                                                                                                    |
| Q01873971     | When using the NNCLI BGP show commands, you should be able to specify a prefix when you use the longer-prefixes option. However, these commands do not allow you to specify the prefix.                                                                                                                           |
| Q01875130     | Several NNCLI BGP commands do not work correctly and peer-group messages of the following types are shown: ERROR: Set Nbr debug: peer-group peer not found. These Privileged EXEC commands include <b>ip bgp stats-clear-counters neighbor &lt;nbr_ipaddr peer-group-name&gt;</b> and <b>ip bgp restart-bgp</b> . |
| Q01754094     | In the NNCLI, you cannot view the SNMP trap log file. The <b>show snmp-server log</b> command is not implemented.                                                                                                                                                                                                 |

| CR references | Description                                                                                                                                                                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01875847     | A warning message is required when the NNCLI flag is changed. When the mode is changed from CLI to NNCLI, the following message should be shown: "Some SNMP commands may not be saved".                                                                                             |
| Q01807816     | In the CLI, the <code>show log file tail name-of-file &lt;filename&gt;</code> command does not produce any output if the <code>tail</code> option is used.<br>Workaround: Redirect the output to another file using the save-to-file option.<br>View the log file in a text editor. |

## BGP resolved issues

Table 19  
BGP resolved issues

| CR references | Description                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------|
| Q01889914-03  | When an AS-List filter is disabled and an IBGP peer is restarted, IBGP will now advertise BGP routes properly. |
| Q01889914-01  | When an AS-List filter is disabled and an IBGP peer is restarted, IBGP will now advertise BGP routes properly. |

## Multicast resolved issues

Table 20  
Multicast resolved issues

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01987813     | Multicast traffic sent with TTL = 1 has been observed to be received by a client in another VLAN when it should not be.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Q01947684     | 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. |
| Q00950235     | IGMPv3 is not backward-compatible with IGMPv1 or IGMPv2 clients. This can have a negative effect on PIM-SSM operations. However, IGMPv3 is not required to use PIM.<br>Workaround: Use IGMPv1 or IGMPv2.                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01891252     | The system has been observed to stop sending MC stream for 10 minutes when the MC sender moves from one port to another.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01679981     | When traffic is flowing in a PIM/SMLT environment and is stopped, active RP entries are not deleted on the MBR neighbor.<br>Release 5.0 does not support MBR functionality over SMLT.                                                                                                                                                                                                                                                                                                                                                                                                                       |

## Quality of Service and filters resolved issues

Table 21

Quality of service and filters resolved issues

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01865206     | When an ACL/ACE filter is applied on a VLAN whose ports belong to different slots, displaying ACL/ACE statistics using the CLI or Device Manager causes high CPU utilization.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Q01450460-01  | <p>When a QoS egress queue template is dynamically applied to a port or a set of ports, egress queue service for some of the ports can stop. The queues become full and those ports no longer transmit packets. Egress queue statistics for the affected ports show dropped pages that increase at the same rate as total pages. Toggling the status of the affected ports from disabled to enabled may restore the ports to normal operation; disable the port before you toggle the administrative status.</p> <p>To avoid this situation, any changes to egress queue templates must be followed by an apply, save config, and then a switch reboot. In a future software version, a reboot will be required for the new egress queue template to take effect, rather than simply performing an apply. Additionally, a warning message for queue template changes will indicate that a reboot is required.</p> |

## Device Manager resolved issues

| CR references | Description                                                                                                                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01875753     | <p>Device Manager incorrectly allows you to add ports from Classic modules to 64-queue egress queue sets.</p> <p>Workaround: Ensure that you only add R or RS module ports to 64-queue egress queue sets.</p> |



---

## Known issues and limitations

---

Use the information in this section to learn more about known issues and limitations. Where appropriate, use the workarounds provided for these.

### Navigation

- [“Release 5.1 known issues” \(page 89\)](#)
- [“Release 5.1 known limitations” \(page 99\)](#)

### Release 5.1 known issues

The following sections list known issues in Ethernet Routing Switch 8600 Release 5.1. These are to be resolved in a future release.

#### Platform known issues

**Table 22**  
**Platform known issues**

| CR references          | Description                                                                                                                                                                                                                                                                                                                                         |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q02025261              | With multiple HA failovers, intermittent connectivity issues may occur.                                                                                                                                                                                                                                                                             |
| Q01993610              | Line RDI is not generated properly as a result of LOS on the 8683XZR module in WAN mode.                                                                                                                                                                                                                                                            |
| Q01993661              | When an RDI-P is received on the XZR module, a "Path RDI" should be shown under the active alarm; however, a "Path AIS" appears.                                                                                                                                                                                                                    |
| Q02005338<br>Q02013534 | With DWDW XFP AA1403001-E5, wavelength is displayed in CLI as 1307nm or 1307.5nm, but the faceplate is labelled as 1310nm.<br><br>The wavelength specifications provided by the XFP manufacturer and reported by the Ethernet Routing Switch 8600 system can vary slightly because of hexadecimal to decimal (and decimal to hexadecimal) rounding. |
| Q01947384              | Ping does not work when the source IP option is set to a circuitless IP interface.                                                                                                                                                                                                                                                                  |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01993806     | <p>Three SONET Alarms are wrongly decoded by the 8683XZR module:</p> <ul style="list-style-type: none"> <li>• Path PLM is interpreted as Path SLM</li> <li>• Path PDI is interpreted as Path SLM Path</li> <li>• TIM is interpreted as No Defect</li> </ul> <p>These three alarms are not supported by Ethernet Routing Switch 8600 Release 5.1.</p>                                                                                                                                                                                                   |
| Q02013161     | <p>With an FPGA upgrade, the PCMCIA is not checked for a file before an error message is displayed. The switch does not search in PCMCIA before displaying the error message, wrongly stating that no such file exists even if the PCMCIA has that file. The switch should search in PCMCIA along with FLASH also. The issue is present for all FPGA upgrade commands.</p>                                                                                                                                                                             |
| Q02002324     | <p>When Autonegotiation is enabled on two switches that are connected to each other using RS modules, and the auto-negotiation-advertisement parameter is set to default on one switch and to 1000-half on the other switch, ping does not work.</p>                                                                                                                                                                                                                                                                                                   |
| Q02004992     | <p>With an HA-enabled switch in synchronized state, if you perform a soft reset on the master CPU (boot), the standby CPU may attempt to become the master. In this case, the master CPU boots and comes up as the standby, and the standby CPU comes up as the master. At this time, the new master shows synchronized state but the standby shows two-way active state, which is wrong.</p>                                                                                                                                                          |
| Q02006487     | <p>When an 8683XZR module in WAN mode receives a LOF, the port correctly detects the LOF, but it does not send out a Line RDI.</p>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Q02006504     | <p>When an 8683XZR module in WAN mode receives a P-PLM (path label mismatch), the alarm raised is path SLM. Path PLM is the SONET term and path SLM is the SDH term. To be consistent, the SONET term should be used when the port is in SONET mode.</p>                                                                                                                                                                                                                                                                                               |
| Q02008247     | <p>With DWDM XFPs, the <b>show system pluggable-optical-modules</b> threshold status is incorrect during transition from a "High Alarm" to a "Low Alarm".</p> <p>When the Rx power is high (beyond the threshold), the threshold status indicator shows "High Alarm", which is correct. However, when the Rx optics is pulled, then the state remains at "High Alarm," even though the indicated power level is -38.200 dBm (which is a "Low Alarm"). The high alarm does not clear in this scenario until the Rx power level goes back to normal.</p> |

| CR references                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01991485                                      | If two Ethernet Routing Switch 8600s are connected using two Ethernet WAN ports using SONET framing and the ports are operationally up, when one of the two ports is changed from SONET framing to SDH framing, no alarm is received. An alarm is needed for the SONET and SDH mismatched connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Q02023608                                      | On 8612XLRS modules with DDM enabled, wait 3 minutes after module initialization before you enter <code>show sys pluggable-optical-modules</code> commands to avoid errors during the initialization.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Q01727720                                      | In the output for the <code>show port error main port &lt;slot/port&gt;</code> command, the <code>FRAMES LONG</code> counter does not increment when the 8648GBRS port receives frames larger than the system MTU size. The same behavior is observed on 8648GTR ports and Gigabit ports on the 8634XGRS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Q02010119                                      | When a dual-CPU switch is configured with a host user ID and password under the <code>/boot/host</code> tree, any attempts to copy files between CPUs fail. To transfer log files to an FTP server, you must configure a user ID and password under <code>/boot/host</code> . However, if you then attempt to copy files from either CPU, the user ID defined under <code>/boot/host</code> is used and the file copy fails.                                                                                                                                                                                                                                                                                                                                                                      |
| Q01811342-02,<br>Q01811343-02,<br>Q01901801-01 | SNMP/JDM connectivity is lost after ERS8600 is reset with <code>hsecure</code> flag is set to true (enabled) and <code>block-snmp</code> flag set to false. To correct, go back and set the <code>block-snmp</code> flag to false. In a future release, <code>hsecure</code> operation will be modified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Q01877552-01                                   | Configuring Distributed MLT on ERS8600 using 8608GBE cards could lead to a ports being put in different STP states in the MLT.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01927867-01                                   | On an ERS8600 running OSPF, the command <code>show vlan info ports &lt;vlan-id&gt;</code> display provides inadequate spacing between the column head VLAN ID and PORT MEMBER when the ports to be displayed are OSPF passive port members.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Q01986390                                      | Force Topology CLIP (Circuitless IP) becomes unconfigured after an HA-CPU failover. Under these considerations, the user must reconfigure the parameter if configured differently than the default value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Q01967211-01                                   | When the remote mirroring destination/termination (RMT) port is used for both local port mirror and remote port mirror, the sniffer can't decode the mirrored traffic and report them as malformed packets. When the RMT port is either used for a remote port mirror or local port mirror, the sniffer can decode the mirrored traffic correctly.<br>When a RMT is configured on an interface, do not configure the same interface as a local "mirroring port". This causes packet corruption on the locally mirrored traffic.<br>In ERS 8600, an interface can be configured as either a local or a remote mirroring destination. The RMT port is exclusively used in Remote mirroring tunnel. The Remote mirroring Destination/Termination port should be disabled to perform local mirroring. |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01750267     | When using High Availability and R mode, if the R mode flag in the configuration files on the two SF/CPU's does not match (or if the primary configuration file is not found on the SF/CPU), the Secondary SF/CPU crashes as reported in Q01832338 and Q01841910.<br>Workaround: Ensure that the configuration files on both the Master and Secondary SF/CPU have the same R mode setting (either enabled or disabled).                                                                      |
| Q01832726     | In a SuperMezz R mode HA-CPU system configured with a dead interval of 3 seconds, when the Master is removed, OSPF neighborship is lost for interfaces configured with low timers (for example, 1 s Hello and 3 s Dead Interval). If failover is triggered by soft-resetting the Master CPU, or the dead interval is 10 s, this issue does not occur.<br>Workaround: Remove the Master CPU during a maintenance window or other low-traffic periods. Or, increase the dead-interval to 10 s. |
| Q01841910     | When using High Availability mode, the Secondary SF/CPU crashes if it cannot find the backup primary configuration file, even if there is a backup configuration file.<br>Workaround: Ensure that the configuration files on both the Master and Secondary SF/CPU have the same R mode setting (either enabled or disabled).                                                                                                                                                                 |
| Q01872749     | If you reset an RS module, a FATAL COP message may be logged. Reset is a maintenance command which should not be used under normal circumstances.<br>Workaround: If you must reset an RS module that is passing traffic, first disable the slot, then reset the slot.                                                                                                                                                                                                                        |

### Switch management known issues

**Table 23**  
**Switch management known issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01924118     | When configuring SSH on the switch, -C and -C2 compression options are accepted, but should be rejected. Subsequent SSH connection are also accepted with no message to the user. The switch should prompt the user with a message stating compression is not supported.                                                                                                                                 |
| Q01576692-01  | With RADIUS accounting of SNMP, when the switch sends start, interim, and stop requests, the packets do not contain the SNMP community string value.                                                                                                                                                                                                                                                     |
| Q01845752-01  | For the Ext-CP Limit feature, 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 references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01972504     | Connection of the Out-of-Band Ethernet Management port to an In-Band I/O port is not recommended, as erroneous behavior on this network, such as a loop, can cause issue with the operation of the SF/CPU module (either 8691 or any 8692 model). The most common issue seen is a loss of file management, and inability to access the /pcmcia directory. A SF/CPU re-boot or reset is needed to clear the condition. The current suggestion to not see such issue is not to connect the OOB port to any In-band I/O port. Such a configuration actually provides no extra value to managing the network. The same functionality can be accomplished, without concern for the above potential situation, by creating a "management" VLAN and assigning a "management" IP address to the VLAN. If OOB management is a desire, then a true OOB management network should be created, and the switches In-Band I/O ports should NOT be part of such a network design. |
| Q01576083     | You can connect through HTTP to an Ethernet Routing Switch 8600 interface that has an IPv4 address, but not to an interface that has an IPv6 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Q01831409     | Using the CLI, when you configure an SNMPv3 target address with an IPv6 address in short format and a port (for example 100::2:162), the software incorrectly assigns part of the IPv6 address as the port number.<br>Workaround: Specify the IPv6 address in long format along with the port number (for example, 1:2:3:4:5:6:7:8:162), or use Device Manager to configure the address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Layer 2 known issues

Table 24  
Layer 2 known issues

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01449886     | The switch can send VLACP PDUs at the short timeout interval when the long timeout interval is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Q01735063     | When the Link Aggregation Control Protocol (LACP) adds a new port to a link aggregation group (LAG), it brings all the ports of the LAG down, which brings the entire interface down. As a result, the multilink trunk is deleted and the VLAN interface is deleted. This causes OSPF to go down.                                                                                                                                                                                                                                          |
| Q01750113     | If you disable RSMLT on a VLAN of a full-mesh triangle SMLT network that carries IPv6 traffic, up to 50% of the IPv6 traffic can be lost.<br>Workaround: Do not disable RSMLT.                                                                                                                                                                                                                                                                                                                                                             |
| Q01850453     | When you convert a single-port MLT to a multiple-port MLT, you cannot set the LACP key for the newly-added ports and form a LAG until you change the key for all ports. You may also receive the following type of error: CPU6 [03/26/08 10:51:50] LACP INFO lacpOperDisablePort: LACP operationally disabled on port 8/23 because the port's capability doesn't match key 64's capability.<br>Workaround: Change all ports to a different key, then change back to the old key. After setting aggregation back to true, the LAG can form. |

## MLT/SMLT known issues

**Table 25**  
**MLT/SMLT known issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01982830     | IPv6 over SMLT or RSMLT is not supported. The following 3 CRs also deal with support of IPv6 over SMLT which is considered an enhancement : Q01834803, Q01756551, 01750113.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Q2016922      | For an IP VPN-lite configuration, where an edge 8600 Cluster is configured to use an SMLT configuration toward the core (most likely square or full-mesh RSMLT), SMLT fast failover cannot always be guaranteed for this portion of the network.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Q01998049     | In RSMLT environments, if an IST peer is powered off for durations longer than the RSMLT holdup-timer, when the powered off peer comes back up on-line, one may see a traffic loss for up to 30 seconds. If this amount of time is unacceptable, it is recommended to perform one of the follow two actions. The recommended action is to configure the RSMLT holdup-timer to infinity on the on-line IST peer, power up the off-line IST peer, and once the SMLT cluster is stabilized, re-configure the RSMLT holdup-timer back to the default setting of 180 seconds. The alternative action is to power up the off-line IST peer during a maintenance windows. |

## Unicast routing known issues

**Table 26**  
**Unicast routing known issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01946521     | In OSPF Router LSA updates, the V-bit is not set, and is always 0.                                                                                                                                                                                                                                                                                                                                                                           |
| Q02008788     | In a square SMLT environment, if OSPF is disabled and reenabled while the IST is down, the OSPF adjacency to one of the non-IST peer boxes may show ExStart state for 5 to 8 minutes. The condition does clear itself in that time frame, and will go to full adjacency.                                                                                                                                                                     |
| Q01976924     | With high route scaling (15 000 routes) and ECMP enabled, error messages can appear on the slave CPU, including:<br>IP ERROR VRF-0: rcIpAddRoute: addIpRoute failed with -102<br>CPU5 [01/07/09 16:04:27] IPMC ERROR ipmSysAddSession FAIL:<br>ipmcSysHashRidAlloc failed S 125.133.180.1 G 224.9.5.0<br>CPU5 [01/07/09 16:04:27] IPMC WARNING ipmcSysHashRidAlloc<br>FAIL: HashRecIndex EOS 8194 S 125.128.168.1 G 224.9.8.0<br>InVlanid 14 |
| Q02010177     | The routing table does not use the preference value specified for a static route if the route has a static ARP entry as the next hop, after disabling and re-enabling the port.                                                                                                                                                                                                                                                              |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q02002237     | <p>At switch boot, IPv6 ND prefixes are created on the switch, but not in the config file. These ND prefixes are created with default valid-life and pref-life parameters. Modifying these values causes the following errors to display on the slave console:</p> <pre>CPU6 [03/05/09 17:53:23] RCIP6 ERROR rcIpv6PrefixTblSetBody: Failed to modify valid life CPU6 [03/05/09 17:53:23] RCIP6 ERROR rcIpv6PrefixTblSetBody: Failed to modify preferred life</pre> <p>These ND prefixes are not saved to the config file after a save config. As a workaround, delete and recreate any IPv6 ND prefixes created upon boot that are not in the config file.</p> |
| Q01872074     | <p>When the switch is in IP VPN Lite mode, the "Assigned Number" portion of the Route Distinguisher (RD) cannot be modified. This occurs for CLI, NNCLI, and Device Manager.</p> <p>Workaround: To change the assigned number, the IP Address and the Assigned Number must be changed at the same time, or the RD must be deleted and re-added (using IP VPN delete).</p>                                                                                                                                                                                                                                                                                       |
| Q01878778     | In BGP update packets, the route origin community type shows the wrong value (0x01 instead of 0x03).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Q01922909-01  | 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, the more specific prefixes do not get redistributed into the OSPF domain. Care must be taken when using such a configuration, to avoid unwanted traffic loss.                                                                                                                                                                                                                                                                                                                                                   |
| Q01464501     | The switch does not send an ICMP Parameter Problem message with an ICMP code value of 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Q01720546     | When an HA failover occurs, it produces BGP errors on the new Master: "CPU5 [07/31/07 11:05:09] BGP INFO PEER_ERROR PKT EVENT(172.16.100.5) Established: remove duplicate peer". No negative effects on functionality occur.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01787988     | For IPv6 traffic, with an IPv6 dest-val of 2301::100 and a src-val of 1301::100, the egress port shown using <code>config sys set hash-calc</code> does not match the original egress port used to send traffic. The port shown by the command is different from the actual port used for traffic forwarding; traffic is forwarded correctly.                                                                                                                                                                                                                                                                                                                   |
| Q01867585     | <p>Configuring a BGP peer group with route-reflection causes configuration anomalies. In particular, you can configure three peers and only two come up; or you can disable route reflection and all peers stay up.</p> <p>Workaround: Do not use BGP route-reflection on any peer groups. BGP peer groups are not supported with route reflection enabled.</p>                                                                                                                                                                                                                                                                                                 |

| CR references | Description                                                                                                                                                                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01878638     | When a VRF with an IP VPN configuration is deleted messages of the following type are displayed on the console: RCMPLS WARNING ipVpnIImAdd line 983: UTAL_LookupArp() failed for NextHop 30.1.30.151 vrfId 100. These messages do not affect functionality.                                               |
| Q01881289     | When a port is untagged and Penultimate Hop Popping (PHP) is set to implicit-null, IP VPN traffic is sent to the wrong queue. The issue occurs on the last-hop Label Switch Router (LSR) (the LSR that sends the service label).<br>Workaround: For untagged ports, set PHP to explicit-null or disabled. |

## Multicast routing known issues

**Table 27**  
**Multicast routing known issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01989326     | MGIDs are not flushed immediately on receiving an IGMPv2 Leave on a snoop-enabled switch. MGIDs are flushed after a delay of 5 to 10 minutes.                                                                                                                                                                                                                                                                                                                                                                                  |
| Q02011440     | SSM channel set to false on receiving traffic                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Q02013130     | Traffic loss is experienced after link failure in PIM-SSM in full mesh topology, which is not supported in Release 5.1.                                                                                                                                                                                                                                                                                                                                                                                                        |
| Q02021453     | RPF checks fail with MSDP peer configured in iBGP configuration.<br>Workaround: Use default Peer or do not use iBGP configuration.                                                                                                                                                                                                                                                                                                                                                                                             |
| Q02003390     | Traffic loss for new streams due to no (S,G) clean up on IST peers. When multiple senders are sending traffic for the same groups (1000 groups) in the same VRF one after another, there may be traffic loss for new streams because this exceeds 2000 (S,G) records in an SMLT/RSMLT environment, due to the (S,G) clean up issue on IST peers.<br>Workaround: Do not send traffic immediately from senders on same groups before timeout or do not send IGMP reports from clients for that particular group for 255 seconds. |
| Q01957808     | Multicast traffic can be observed going to switches in a full-mesh or square RSMLT/SMLT configuration, even though no ports on those switches are members of the multicast group. This is per design for the switch CPP to know where the multicast groups exist, and is normal expected behavior. This group traffic will be dropped at the CPP after learning and not forwarded.                                                                                                                                             |
| Q01891238     | Multicast streams that move from ingressing an MLT to ingressing a different port may experience interruption.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Q01866720     | The output of the <code>show ip pim interface</code> command indicates that the operational status (OPSTAT) is down even if the PIM interface is enabled and active. This does not impact traffic.                                                                                                                                                                                                                                                                                                                             |

## CLI and NNCLI known issues

Table 28  
CLI and NNCLI known issues

| CR references          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q02000473              | After enabling Hsecure on the switch and saving the configuration, the CLI prompt should not be returned to the user until the configuration save is complete. Currently, the switch displays the following error: Another show or save in progress. Please try the command later.                                                                                                                                                                                          |
| Q02021276              | Configuration information for PIM virtual neighbours in VRF are saved under multiple headers. The general practice for saving configurations in config.cfg is to save configurations for all the VRFs under a single header. The configurations for PIM virtual neighbour do not follow the convention as a new header is created for each VRF.                                                                                                                             |
| Q01467778<br>Q02010136 | The System Messaging Platform (SMP) command <b>config log transferFile &lt;ID&gt; filename &lt;str&gt;</b> , used to modify the default file name, does not function properly. The optional filename string does not work. Instead, the transferFile name is always the standard SMP log file name of xxxxxx.sss, where xxxxxx is the last three bytes of the chassis base MAC address, and sss is the SF/CPU slot number. Nortel recommends that you not use this command. |
| Q01948346              | The <b>copy</b> command does not work properly with FTP debug turned on.                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q02010142              | When sending traps to an Element Manager, the switch only uses the IP address specified by the first entered <b>sender-ip &lt;dest-ip&gt; &lt;source-ip&gt;</b> command. It is possible to specify multiple sender IPs and each should use a different IP as specified in this command.<br>The switch uses the IP address of the physical VLAN of the first entry in the target-address table.                                                                              |
| Q01623093              | In the CLI, when you enter the <b>config/ip/static-route# create &lt;ipaddr/mask&gt; next-hop ?</b> command, only the <b>vrf &lt;value&gt;</b> parameter is shown in response. In spite of this, you can still configure all parameters:<br><b>config/ip/static-route# create &lt;ipaddr/mask&gt; next-hop &lt;ipaddr&gt; cost &lt;value&gt; [preference &lt;value&gt;] [local-next-hop &lt;value&gt;] [next-hop-vrf &lt;value&gt;]</b>                                     |
| Q01794612              | Entering the CLI <b>reset</b> command for an SF/CPU module with SuperMezz should implement a cold boot on both the SF/CPU and SuperMezz, but instead incorrectly performs a warm boot on the SuperMezz and correctly performs a cold boot on the SF/CPU.                                                                                                                                                                                                                    |

## Quality of Service and filters known issues

Table 29  
QoS and filters known issues

| CR references | Description                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------|
| Q01904368     | Unable to mirror IPv6 packets received on a port when the ACE filter action on that port is set to deny. |

| CR references | Description                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01917455     | On RS modules, if you enable access-diffserv ( <b>enable-diffserv true</b> ) on ingress or egress ports, IPv6 traffic on the ports is dropped and the IPv6 neighbor state changes to INCOMPLETE. |
| Q01793851     | RS modules do not support egress flow-based mirroring for IP VPN Lite traffic.                                                                                                                   |
| Q01814530-02  | ACE filters applied to IST ports may malfunction after the switch reboots. Workaround: Disable and enable the malfunctioning ACE.                                                                |
| Q01868828     | If you enable mirroring on an ACL-based filter whose actions re-mark the DSCP bit, the DSCP bit and QoS egress queues are incorrectly assigned.                                                  |

### Device Manager known issues

**Table 30**  
**Device Manager known issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q02006489     | No help for DDI/SFP tab in Device Manager. On clicking help in the DDI/SFP tab, the following error appears:<br>Unable to open "D:\Program Files\Nortel\JDM617B03\helpaccelar2k/v510/rcPlugOptModPort.html"                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Q01984969     | When doing an SNMP walk using Device Manager for any VLAN, the index of MIB objects is out of sequence for vlan fdb entry object. The warning message "The index of MIB objects is out of sequence!" is displayed in the Device Manager log.                                                                                                                                                                                                                                                                                                                                                                                              |
| Q01995139     | Device Manager and agent support are not available for PIM Virtual interfaces. The Device Manager equivalent for these commands is not present: <ul style="list-style-type: none"> <li><b>config ip pim interface &lt;ipaddr&gt; virtual-neighbor add &lt;ipaddr&gt;</b></li> <li><b>config ip pim interface &lt;ipaddr&gt; virtual-neighbor delete &lt;ipaddr&gt;</b></li> <li><b>config ip pim interface &lt;ipaddr&gt; virtual-neighbor info</b></li> </ul>                                                                                                                                                                            |
| Q01641989-01  | In Device Manager, under the <b>IP, IGMP, StreamLimit</b> tab, you cannot configure the MaxNumOfStreamLimit variable unless you first enable the StreamLimitEnable option. However, if the number of groups joined on an interface already exceeds the configured MaxNumOfStreamLimit, you cannot enable StreamLimitEnable.<br>In this case, to modify the MaxNumOfStreamLimit for an interface, you can use the <b>VLAN-&gt;VLANs-&gt;IP-&gt;IGMP</b> tab. With this tab, it is possible to set the variable regardless of the state of the StreamLimitEnable option. Otherwise, you can use the CLI to configure the IGMP stream limit. |

| CR references | Description                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01461573     | In Device Manager, the Clear Counters button is missing for Graph > Port > Egress Queue Set.                                                                                                                                                            |
| Q01883183     | When using Device Manager, IP VPN routes in a VRF are not flushed when the import route target (RT) is deleted.<br>Workaround: To delete the routes, before you delete the RT, in the IP-VPN, VPN tab, ensure that the import RT field is not selected. |

## Release 5.1 known limitations

Use the information in this section to learn more about known limitations. These CRs are classified as operation not to be changed.

### Platform limitations

**Table 31**  
**Platform limitations**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01279047-01  | Disabling Auto-negotiation on 1000BASE-T ports with speed of 1000 Mbps can result in inconsistent behavior. Auto-negotiation cannot be disabled on the 8648GTR module, but can be disabled on an 8630GBR with an SFP.<br>Workaround: Nortel recommends that you enable Auto-negotiation on all 1000BASE-T ports when they operate at 1000 Mbps.                                                                                                                                                      |
| Q01347146-01  | 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.<br>Workaround: Disable SFFD and instead use VLACP.                                                                                                                                                                                                                                                 |
| Q01464922     | An 8630GBR module using Single Fiber Fault Detection (SFFD) connected to an Ethernet Switch 470 does not recognize a single-fiber fault. SFFD is not supported when an Ethernet Routing Switch 8600 is connected to an Ethernet Switch 470.<br>Workaround: Disable SFFD and instead use VLACP.                                                                                                                                                                                                       |
| Q01971595     | Pasting configurations, either via console or telnet/SSH, is not recommended and if performed should only be done one line at a time, versus larger text files. This is independent of any control the user may have over how 'slow' their actual input maybe. The reason for is that neither the console or telnet/SSH functions are high priority, and the switch could be servicing some other task and therefore miss lines in the config paste. Optionally, use TFTP/FTP and 'source' the file. |
| Q01938504     | 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.                      |

| CR references              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01871916                  | Broadcast and multicast rate limiting may occasionally cause packets to be dropped when traffic levels are below the configured rate limit when configured on classic modules.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Q01774929                  | The event of discontinuing logging to PCMCIA by the executing the CLI command <code>pcmcia-stop</code> is not logged into the log file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Q01848350                  | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Q01937023                  | A “save config” executed from the standby CPU will result in the <code>pcap.cfg</code> data replacing the configuration data stored in <code>config.cfg</code> on the standby CPU. Changes to the standby CPU configuration should ONLY be executed from the master CPU, via either “save config” ( <code>savetostandby</code> flag set true or “save config standby <config file name>”.                                                                                                                                                                                                                                                                 |
| Q01926665-02               | SLT configurations where the SLT port and all the IST_MLT ports are configured on the same physical module (non-best practice design; IST_MLT should always be a D-MLT), the switch can experience FDB entry learning issues for MAC addresses learned across the IST; this can lead to 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. If seen and to resolve, FDB entries will be relearned after the <code>fdb-ageout</code> time, or can be manually relearned by performing an <code>fdb-entry flush</code> on the relevant VLANs. |
| Q01904966,<br>Q01869054-01 | On an ERS8600 running in SMLT environment using classic (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.                                                                                                                                                                                                                                                                                                                                                                                                              |
| Q01932414                  | 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.                                                                                                                                                                                                                                                                                                                                                                                                       |
| Q01952948                  | 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.                                                                                                                                                                                                                                                       |
| Q01880938-01               | LACP short timers for RSMLT/SMLT configurations must be set to 550 milliseconds or higher.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Q01885138-04               | Traffic loss has been observed for 1-2 seconds after a port link up occurs. This is caused by the edge box transmitting traffic immediately upon connecting to the aggregation (IST core) box. As soon as the link is enabled, it takes a while before VLACP engages and stops the traffic from flowing on that link.                                                                                                                                                                                                                                                                                                                                     |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01986665-01  | <p>The PCMCIA card may become inaccessible and the following error may be shown:</p> <p><b>0x53e7ab0 (tLoggerTask): disk cache error: device fe21488 block 57 errnoc0003, disk write failed, data loss may have occurred.</b></p> <p>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.</p>                                                                                                                                                                                                                      |
| Q01986615     | <p>If you have a IPVPN Lite route, with multiple IGP nexthops towards the MP-BGP network (by virtue of ECMP being enabled in VRF 0) and ECMP is also enabled in a non-zero VRF, then the VRF route is installed multiple times in Route Table Manager (RTM). The datapath will therefore use each IGP nexthop, as expected. Now, if the user disables ECMP in the VRF, the additional ECMP routes from the non-zero VRF are deleted but the IPVPN lite route is not removed. If then some network failover happens, under these circumstances, the system incorrectly thinks that the underlying route is still available. This could cause unwanted traffic interruption.</p> <p>Work around: Restart IPVPN if this event occurs.</p> |
| Q01972028     | <p>When operating in HA mode, if the Standby and Master SF/CPU's are rebooted quickly in succession within a few seconds in that order, a race condition can occur in which R/RS line cards do not get rebooted as part of the Standby and Master reboot process. Instead they get rebooted in a delayed fashion during the new Master SF/CPU initialization. However, they would function normally after that. To avoid this delayed reboot condition and an unnecessary alarm, Nortel recommends to not reboot both the Standby and Master SF/CPU's in quick succession or simultaneously. By leaving a gap of around 10 seconds, this problem could be avoided.</p>                                                                 |
| Q01994125     | <p>When connecting ports out of order (Port 1/1 to Port 1/10 for example) on two boxes, some LACP links may go to standby instead of active.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Q01986915     | <p>MLT operation for ingress R/RS ports is a L4 (TCP/UDP port) hash, with the exception of fragmented data streams, as the fragmented packets have no L4 header. Therefore the hash is L3 based, which could lead to potential issues with MLT or SMLT connected devices (especially state full firewalls) that want all streams to be received on the same physical link of the MLT. A quick test for such conditions is to force the MLT to become one port by disabling the other ports in the MLT. If operation now improves, this is probably the situation. Nortel recommends a change in the physical connectivity of such devices, until this situation can be improved in a future software release.</p>                      |
| Q01443107-01  | <p>When you initialize an ATM card without a proper .ldd file, you cannot gain access to the switch through the console, Telnet, or Device Manager.</p> <p>Workaround: Do not boot a system with an ATM module installed without the proper image file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01450334-01  | The switch does not drop packets with unknown source MAC addresses when the number of learned fdb-entries reaches the maximum limit. The opposite operation is documented incorrectly in the user documentation.                                                                                                                                                                                                                                                                                                           |
| Q01470456     | When the Packet Capture Tool (PCAP) buffer fills up, a Standby SF/CPU exception can occur. The exception only occurs if you use a console connection to the Standby SF/CPU. This only occurs when the autosave feature is enabled and is set to transfer to a network device. The issue does not occur when you save to a PCMCIA card.                                                                                                                                                                                     |
| Q01450691     | 8616GT/GTE modules discard jumbo frames by hardware design. The 8616GTE or 8616SXE modules do not support jumbo frames.                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01450708     | When you create an MLT with more than one port that ends in an 8648GTR module, and the MLT is connected to a 10/100 Mbps device with Auto-negotiation enabled, the link state for the first port is 100 Mbps full duplex and state of all other MLT ports is 10 Mbps half duplex.<br>Workaround: Adjust the speed and duplex modes manually on a per-port basis.                                                                                                                                                           |
| Q01848971     | Ping and traceroute may fail for VRF, IP VPN, or MPLS routes if large packet sizes are used for the operation ().<br>Workaround: do not use packet sizes larger than the following: <ul style="list-style-type: none"> <li>• Ping for VRF Lite: 1480 bytes</li> <li>• Traceroute for VRF Lite: 1444 bytes</li> <li>• Ping for IP VPN with MPLS: 1480 bytes</li> <li>• Traceroute for IP VPN with MPLS: 1444 bytes</li> <li>• Ping for IP VPN Lite: 1446 bytes</li> <li>• Traceroute for IP VPN Lite: 1444 bytes</li> </ul> |
| Q01857099     | To configure a WSM in NNCLI mode, you must load the configuration file in NNCLI mode. However, this causes errors.<br>Workaround: Use the CLI or Device Manager to configure the WSM.                                                                                                                                                                                                                                                                                                                                      |
| Q01644085     | When a copper GBIC is present in an 8630GBR module port with Auto-Negotiation disabled, the port is always in an up state, even if no Ethernet cable is plugged in.<br>Workaround: Auto-Negotiation must be enabled for GBIC ports running 1 Gigabit Ethernet. Keep Auto-Negotiation enabled.                                                                                                                                                                                                                              |

## Switch management limitations

**Table 32**  
**Switch management limitations**

| CR references | Description                                                                            |
|---------------|----------------------------------------------------------------------------------------|
| Q01393349-01  | On the ERS8600 switch, the -f option cannot be used with SSH non interactive commands. |

| CR references           | Description                                                                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01380593,<br>Q01194685 | The switch does not add standby ports to a link aggregation group (LAG) properly. Currently, the Ethernet Routing Switch 8600 supports a maximum of eight ports, active or standby, in any one LAG. |
| Q01436111               | The switch does not display all SNMPv3 target parameters when you enter the <code>show config</code> command. This is the expected behavior.                                                        |
| Q01454269               | The switch does not function when two Secure Shell (SSH) sessions originate from the same PC port.                                                                                                  |

## Layer 2 limitations

**Table 33**  
**Layer 2 limitations**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01449886     | The switch can send VLACP PDUs at the short timeout interval, even though the long timeout interval is configured. This operation has no negative affect on VLACP operation, to other ERS 8600s, or to any other Nortel ERS switches running VLACP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Q01352932     | In a pure Layer 2 point-to-point configuration, using VLACP fast timers, 100 ms convergence is not guaranteed when ECMP is implemented on top of the link. Keep the configuration as simple as possible, with routed VLANs, but without ECMP. For fast convergence, Nortel recommends that you use RSMILT rather than ECMP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Q01422264     | <p>In a triangle setup, if an Ethernet Routing Switch 8600 boots up with Rapid Spanning Tree Protocol (RSTP) mode, and has a link that includes a VRRP neighbor, the VRRP state goes directly to MASTER instead of waiting for the VRRP hold-down timer to expire. This is due the fact that, in some cases, all the VLAN ports transition to a blocking state and RSTP may bring down the VLAN and IP interface.</p> <p>This is a timing issue that can occur in mixed RSTP/VRRP environments, and occurs because RSTP needs time to converge. The result is that a loss of connectivity occurs until VRRP converges.</p> <p>Workaround: To avoid this situation, add additional active port members to the VLAN. Then, the VLAN does not go down, and the loss of connectivity does not occur.</p> |
| Q01453004     | With HA enabled, the switch does not flush the VR ARP entry from the ARP table when you delete the VR ID from both switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Q01454547     | You can configure Link Aggregation/MLT group IDs from 1 through 256. The switch supports a total of 128 MLT groups. Nortel suggest that users start numbering S-SMLT or SLT IDs to avoid overlapping MLT IDs in the 1 through 128 range. Note: SLT IDs range from 1 through 512. If you start at 129, a switch can support the full capacity of 383 SLT ports or IDs.                                                                                                                                                                                                                                                                                                                                                                                                                                |

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01577517-01  | The 8692 SF/CPU module RSP memory has a limitation that can cause an out-of-memory issue with source MAC-based VLANs if a large number of MAC addresses (4000+) and ports (25+) are used.<br>Workaround: limit the number of MAC addresses and ports for source MAC address-based VLANs.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Q01463337     | When you delete a large number of VLANs, Cop-SW errors may appear: "CPU6 [09/25/06 15:23:06] COP-SW ERROR Slot 7: ercdFilterVlanRadixDelete ERROR: DELETING Vlan=318 from FILTER VLAN RADIX LOOKUP Table". These errors do not have any negative affect on functionality, because everything is deleted when the script finishes.                                                                                                                                                                                                                                                                                                                                                                                                 |
| Q01810668     | When you install an RS module and have SMLT and LACP configured, LACP goes down, and the switch shows the message "swA:6# CPU6 [01/09/08 14:39:28] LACP INFO lacpOperDisablePort: LACP operationally disabled on port x/y because the port's capability doesn't match key z's capability".<br>Workaround: Globally disable and enable LACP to bring the ports back to the MLT. Further, for reliable LACP operation, use CANA, or, if CANA is not supported, enforce a speed and duplex mode on all 10/100/1000 Mbps and 10/100 Mbps ports. See <a href="#">"Mixed-module MLT/LAG support for Release 5.0 and later" (page 39)</a> .                                                                                              |
| Q01885788     | After an upgrade of an IST peer, you may see messages of the following type on the console: CPU6 [06/05/08 20:15:18] IPMC WARNING ipmSysModifySessionParams ipmSysArModifyIpmcRec FAIL.<br>Workaround: After an IST peer is upgraded and restarted, wait until the entire system is stable prior to upgrading the other IST peer.<br>Stabilization time depends on the complexity and size of the network (for example, the number of MAC and ARP records, routes, and the protocols used). Wait for the Layer 3 protocols, especially multicast protocols, to settle before you restart the other peer. If Layer 3 protocols are not in use, wait until the FDB and ARP tables on both peers report a similar number of entries. |

### MLT/SMLT limitations

**Table 34**  
**MLT/SMLT limitations**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01971344     | In the case of a full-mesh SMLT configuration between 2 Clusters running OSPF (more likely an RSMLT configuration) because of the way that MLTs work in regards to CP generated traffic, it is highly recommended that the MLT port (or ports) that form the square leg of the mesh (versus the cross connect) be placed on lowered number slot/port, than the cross connections. The reason for this is because CP generated traffic is always sent out on the lowered numbered ports when active. Using this recommendation will keep some OSPF adjacency up if all the links of the IST fail. Otherwise the switches which have a failed IST could lose complete OSPF adjacency to both switches in the other Cluster and therefore become isolated. |

## Unicast routing limitations

**Table 35**  
**Unicast routing limitations**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                      |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01845219-01  | 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. |
| Q01440964-01  | It may appear that the <code>show ip ospf lsdb</code> command does not completely execute. When the link state database (LSDB) synchronizes due to a change in the network, the command cannot return all the requested data.<br>Workaround: after the switch completes the LSDB synchronization process, run the command again.                                 |
| Q01449035     | IPv6 addresses remain in the routing table after you disable the VLAN.                                                                                                                                                                                                                                                                                           |
| Q01835057     | To maintain backward-compatibility with pre-5.0 releases, if the management IP address network mask is equal to the natural network mask, then the mask is not displayed in response to the <code>config bootconfig net mgmt info</code> command. If the configured network mask is not equal, it will be shown. This is the correct behavior.                   |

## Multicast routing limitations

**Table 36**  
**Multicast issues**

| CR references | Description                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01437712     | The <code>show ip pim mroute</code> command may display an incorrect incoming port after the corresponding SMLT port is disabled and re-enabled.          |
| Q01439217-01  | The switch sends no ICMP message when it receives IPv6 packets destined to an IPv6 multicast address with a destination header option type equal to 0x80. |

## QoS and filters limitations

**Table 37**  
**QoS and filters issues**

| CR references | Description                                                                                  |
|---------------|----------------------------------------------------------------------------------------------|
| Q01443034     | The CLI allows you to configure mutually-exclusive Access Control Template (ACT) properties. |
| Q01450812     | The <code>show qos stats policy all</code> command does not display correct statistics.      |

## Device Manager limitations

**Table 38**  
**Device Manager issues**

| CR references | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q01374909     | <p>Device Manager does not allow all fields to be configured for PING and Trace Route. This is as designed.</p> <p>Trace Route does not allow the following fields:</p> <ul style="list-style-type: none"><li>• TargetAddressType</li><li>• DSField</li><li>• SourceAddressType</li><li>• IfIndex</li><li>• MiscOptions</li><li>• StorageType</li><li>• CreateHopsEntries</li><li>• Type</li></ul> <p>PING does not allow the following fields:</p> <ul style="list-style-type: none"><li>• TargetAddressType</li><li>• StorageType</li><li>• Type</li><li>• SourceAddressType</li><li>• IfIndex</li><li>• DSField</li></ul> |
| Q01444776     | Device Manager provides no option to configure the IPv6 retransmit-timer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Q01444777     | Device Manager provides no option to configure the IPv6 reachable-time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Q01445777     | The Device Manager scroll bars on the ACL-Insert ACL-ActID interface do not function properly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## MIB limitations

**Table 39**  
**MIB issues**

| CR references | Description                                                                                                        |
|---------------|--------------------------------------------------------------------------------------------------------------------|
| Q01274812-01  | The description for rcStgTaggedBpduVlanId is incorrect. The range -2147483648 to 2147483647 should read 1 to 1024. |

## Customer service

---

Visit the Nortel Web site to access the complete range of services and support that Nortel provides. Go to [www.nortel.com](http://www.nortel.com), or go to one of the pages listed in the following sections.

### Navigation

- “Updated versions of documentation” (page 107)
- “Getting help” (page 107)
- “Express Routing Codes” (page 107)
- “Additional information” (page 108)

### Updated versions of documentation

You can download and print the latest versions of Nortel Ethernet Routing Switch 8600 NTPs and Release Notes directly from the Internet at [www.nortel.com/documentation](http://www.nortel.com/documentation).

### Getting help

If you purchased a service contract for your Nortel product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel service program, you can get help by contacting one of the Nortel Technical Solutions Centers found at [www.nortel.com/callus](http://www.nortel.com/callus); or visit our Technical Support site at [www.nortel.com/support](http://www.nortel.com/support).

### Express Routing Codes

An Express Routing Code (ERC) is available for many Nortel products and services.

When you use an ERC, your call is routed to a technical support person who specializes in supporting that particular product or service. To locate an ERC for a product or service, go to [www.nortel.com/erc](http://www.nortel.com/erc).

## Additional information

Use the information in the following table to access other areas of the Nortel Web site.

| For information about            | Contact                                                                              |
|----------------------------------|--------------------------------------------------------------------------------------|
| Contact Us                       | <a href="http://www.nortel.com/contactus">www.nortel.com/contactus</a>               |
| Documentation feedback           | <a href="http://www.nortel.com/documentfeedback">www.nortel.com/documentfeedback</a> |
| Products (marketing)             | <a href="http://www.nortel.com/products">www.nortel.com/products</a>                 |
| Partner Information Center (PIC) | <a href="http://www.nortel.com/pic">www.nortel.com/pic</a>                           |
| Register                         | <a href="http://www.nortel.com/register">www.nortel.com/register</a>                 |
| Search                           | <a href="http://www.nortel.com/search">www.nortel.com/search</a>                     |
| Services                         | <a href="http://www.nortel.com/services">www.nortel.com/services</a>                 |
| Training                         | <a href="http://www.nortel.com/training">www.nortel.com/training</a>                 |

---

# Index

---

## A

Advanced filters  
guidelines 75

## B

boot source  
SuperMezz 69

## C

chassis  
High Performance 65  
compatibility  
module and chassis 64  
module and component 58  
module and feature 63  
module and mode 63  
software and hardware 57

## D

Device Manager  
and Linux 77  
and Solaris 77  
and Windows 77  
installation requirements 76  
DOSFS 41  
DVMRP considerations 71

## E

ESM considerations 75

## F

file names 32  
Fixes from previous releases 40

## H

HA mode  
considerations 73

## I

I/O module considerations 39  
IGMPv3 18  
IST and VLACP considerations 71

## K

known issues  
CLI and NNCLI 97  
Device Manager 98  
Layer 2 93  
MLT/SMLT 94  
multicast routing 96  
platform 89  
QoS and filters 97  
switch management 92  
unicast routing 94

## L

license  
temporary 75  
limitations  
Device Manager 106  
Layer 2 103  
MIB 106  
MLT/SMLT 104  
multicast routing 105  
platform 99  
QoS and filters 105  
switch management 102  
unicast routing 105  
loop prevention 66

**M**

- MIB 76
- Mixed-module MLT/LAG support 39
- MLT/LACP
  - and port speed 40
- MPLS
  - and MTU 70
- MSDP 18
- MSDP considerations 73
- multicast on a VRF 18
- multicast static IP route table 18

**N**

- nonsupported modules 39
- notifications 76

**P**

- PCAP considerations 74
- performance
  - module and chassis 64
- power and cooling management 43

**R**

- resolved issues
  - BGP 86
  - CLI and NNCLI 85
  - Device Manager 87
  - IP unicast 83–84
  - layer 2 switching 79
  - multicast 86
  - platform 81
  - QoS and filters 87
  - switch management 80
- RFC 76

**S**

- scaling 52
- SF/CPU
  - protection and loop prevention 66
- SMLT
  - and switch clustering 66
- SMLT considerations 71
- SNMP considerations 70
- Software licensing 31
- Software Release 3.7 support 77
- SuperMezz 36

**T**

- traps 76

**U**

- upgrade considerations 41
  - changes to procedures 46
  - IST 45
  - power management 42
  - pre-5.0.1.0 software 47
- upgrade paths 41

**V**

- VLACP and IST considerations 71

**W**

- WSM considerations 75



Nortel Ethernet Routing Switch 8600

## Release Notes — Software Release 5.1

Release: 5.1

Publication: NN46205-402

Document revision: 03.02

Document release date: 4 June 2009

Copyright © 2008-2009 Nortel Networks. All Rights Reserved.

While the information in this document is believed to be accurate and reliable, except as otherwise expressly agreed to in writing NORTEL PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED. The information and/or products described in this document are subject to change without notice.

Nortel, Nortel Networks, the Nortel logo, and the Globemark are trademarks of Nortel Networks.

All other trademarks are the property of their respective owners.

To provide feedback or to report a problem in this document, go to [www.nortel.com/documentfeedback](http://www.nortel.com/documentfeedback).

[www.nortel.com](http://www.nortel.com)

