



Release Notes — Software Release 7.1

Avaya Ethernet Routing Switch 8800/8600

7.1
NN46205-402, 05.01
February 2011

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on its Hardware and Software (“Product(s)”). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support Web site: <http://support.avaya.com>. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](http://support.avaya.com/licenseinfo/) ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or Hardware provided by Avaya. All content on this site, the documentation and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third-party components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements (“Third Party Components”), which may contain terms that expand or limit rights to use certain portions of the Product (“Third Party Terms”). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://support.avaya.com/Copyright>.

Preventing Toll Fraud

“Toll fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud Intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support Web site: <http://support.avaya.com>. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners, and “Linux” is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: <http://support.avaya.com>.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://support.avaya.com>.

Contents

| | |
|--|-----------|
| Chapter 1: Introduction | 7 |
| Navigation | 7 |
| Chapter 2: New features in this release | 9 |
| SPBM | 9 |
| SPBM MGID usage | 10 |
| CFM | 10 |
| 8800 series I/O modules | 11 |
| SMLT interoperability with MSTP | 11 |
| 4-byte AS support for BGP | 12 |
| HTTPS support for EDM | 12 |
| RADIUS support for EDM | 12 |
| EDM VRF context switching enhancement | 12 |
| EDM navigation enhancement | 13 |
| Multiple port selection and monitoring in EDM | 13 |
| VRF name increase from 32 to 64 characters | 13 |
| BFD enhancements | 14 |
| SLPP enhancements | 14 |
| RSP Trace Filtering | 15 |
| show command timestamp | 15 |
| CP limit statistics | 15 |
| CLI wrapper for shell debug commands | 15 |
| Flash file system enhancements | 16 |
| Chapter 3: Important notices | 17 |
| Navigation | 17 |
| Supported hardware and software compatibility | 18 |
| Unsupported hardware for Release 7.1 | 21 |
| Supported software and hardware scaling capabilities | 22 |
| Software licensing | 26 |
| File names for this release | 27 |
| Important information and restrictions | 29 |
| Fixes from previous releases | 30 |
| Important information and restrictions navigation | 30 |
| SuperMezz, SF/CPU memory, and upgrades | 30 |
| Compact flash card display on 8895 SF/CPU | 31 |
| Proper care of external compact flash and PCMCIA cards | 31 |
| Pasting configurations into the configuration file | 31 |
| EDM considerations | 32 |
| I/O module considerations | 36 |
| MLT/LAG considerations | 36 |
| Console connection considerations | 36 |
| DHCP snooping considerations | 36 |
| Supported upgrade paths | 37 |
| General upgrade considerations | 37 |
| Upgrade considerations for Release 7.1 | 37 |
| Upgrading from 8692 SF/CPU with SuperMezz to 8895 CPU | 38 |
| Upgrade considerations: DOSFS with upgrades from pre-Release 5.0 | 39 |

| | |
|--|-----------|
| Upgrade considerations: Power Management..... | 40 |
| Disabling power and cooling management..... | 41 |
| Upgrade considerations: IST..... | 43 |
| Pre-release 5.1 upgrades considerations: specifying license file location..... | 43 |
| Considerations for upgrades from 5.0-based code releases..... | 44 |
| Configuration file modifications for BGP upgrades from release 4.x code..... | 44 |
| SMLT switch cluster upgrade considerations..... | 45 |
| High Availability mode considerations..... | 47 |
| Ongoing considerations..... | 48 |
| Module and chassis compatibility and performance considerations..... | 48 |
| High Performance chassis..... | 49 |
| Switch clustering topologies and interoperability with other products..... | 50 |
| SF/CPU protection and loop prevention compatibility..... | 50 |
| Switch behavior during boot cycle and redundant configuration files..... | 50 |
| Configuring primary, secondary, and tertiary boot sources..... | 52 |
| OSPF warning message..... | 53 |
| MPLS considerations..... | 54 |
| IPv6 considerations..... | 54 |
| SNMP considerations..... | 54 |
| DVMRP considerations..... | 55 |
| SMLT considerations..... | 55 |
| RSMLT considerations..... | 55 |
| IST considerations..... | 56 |
| 60 day trial license..... | 56 |
| Advanced filter guidelines..... | 57 |
| MTBF for 1 Gig SFPs..... | 57 |
| Supported standards, RFCs, and MIBs..... | 58 |
| Supported traps and notifications..... | 58 |
| Chapter 4: Resolved issues..... | 59 |
| Platform resolved issues..... | 59 |
| Switch management resolved issues..... | 60 |
| KHI resolved issues..... | 60 |
| Layer 2 resolved issues..... | 61 |
| MLT/SMLT resolved issues..... | 61 |
| Unicast routing resolved issues..... | 61 |
| Multicast routing resolved issues..... | 62 |
| IPv6 resolved issues..... | 62 |
| CLI and ACLI resolved issues..... | 63 |
| Enterprise Device Manager resolved issues..... | 64 |
| Off-box EDM plug-in resolved issues..... | 66 |
| Chapter 5: Known issues and limitations..... | 67 |
| Navigation..... | 67 |
| Release 7.1 known issues..... | 67 |
| Previously reported known issues..... | 69 |
| Platform known issues..... | 69 |
| Switch management known issues..... | 71 |
| KHI known issues..... | 72 |
| Layer 2 known issues..... | 72 |
| MLT/SMLT known issues..... | 73 |
| Unicast routing known issues..... | 73 |

| | |
|--|-----------|
| Multicast routing known issues..... | 74 |
| IPv6 known issues..... | 74 |
| CLI and ACLI known issues..... | 75 |
| Enterprise Device Manager known issues..... | 75 |
| Chapter 6: Customer service..... | 79 |
| Navigation..... | 79 |
| Getting technical documentation..... | 79 |
| Getting product training..... | 79 |
| Getting help from a distributor or reseller..... | 79 |
| Getting technical support from the Avaya Web site..... | 80 |

Chapter 1: Introduction

This document describes important notices and fixed and known issues for Avaya Ethernet Routing Switch 8800/8600 release 7.1 software.

Ethernet Routing Switch 8800/8600 release 7.1 supports the 8895 Switch Fabric/CPU Module. When an 8000 Series Chassis is equipped with the 8895 SF/CPU, this system is known as an Ethernet Routing Switch 8800; conversely, when equipped with an 8692 SF/CPU module (with SuperMezz), the system is known as an Ethernet Routing Switch 8600. Ethernet Routing Switch 8800/8600 release 7.1 software can only operate on an Ethernet Routing Switch 8800/8600 system with appropriate hardware configurations.

Refer to the following sections of the Release Notes for additional detailed information regarding the supported ([Supported hardware and software compatibility](#) on page 18) and unsupported ([Unsupported hardware for Release 7.1](#) on page 21) combinations of hardware and software, as well as new feature descriptions.

Navigation

- [New features in this release](#) on page 9
- [Important notices](#) on page 17
- [Resolved issues](#) on page 59
- [Known issues and limitations](#) on page 67
- [Customer service](#) on page 79

Chapter 2: New features in this release

The following sections describe the new features for the Avaya Ethernet Routing Switch 8800/8600 Release 7.1.

SPBM

Release 7.1 of the Ethernet Routing Switch 8800/8600 supports the IEEE 802.1aq standard of Shortest Path Bridging MACinMAC (SPBM). SPBM makes network virtualization much easier to deploy within the Enterprise environment, reducing the complexity of the network while at the same time providing greater scalability.

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core control plane to a single protocol which can provide virtualization services for both layer 2 and layer 3, on a common Ethernet infrastructure using a pure Ethernet technology base. SPBM allows for layering the Ethernet network into edge and core domains with complete isolation between their MAC addresses. This technology provides all the features and benefits required by Carrier-grade deployments to the Enterprise market without the complexity of alternative technologies traditionally used in Carrier deployments (typically MPLS). SPBM integrates into a single control plane all the functions that MPLS requires multiple layers and protocols to support.

SPBM provides any-to-any connectivity in a network in an optimized, loop-free manner. It employs shortest-path trees to each destination, without the long convergence delays experienced with Spanning Tree Protocol. To do that, SPBM uses Intermediate System to Intermediate System (IS-IS) link state routing protocol to learn and distribute network information. IS-IS dynamically learns the topology of a network and uses its inherent knowledge to construct shortest path unicast and multicast trees from every node to every other node in the network. Also, unlike Spanning Tree Protocol, IS-IS does not block ports to provide a loop free topology, so bandwidth is not wasted.

The SPBM components introduced in this release are:

1. Shortest Path Bridging (IEEE 802.1aq) for simple and safe **VLAN extensions** across a network. SPBM does not use spanning tree, and all its links are active.
2. SMLT for dual-homing of non-SPBM switches to a pair of SPB/IST switches.
3. SPBM/IP for simple and safe **VRF extensions** across a network infrastructure without OSPF or BGP.
4. InterISID routing for routing of L2 VPNs in the SPB domain

For more information about SPBM, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Shortest Path Bridging MAC (SPBM)* (NN46205–525).

SPBM MGID usage

The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the data is directed to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs. The system also reserves a small number of MGIDs.

SPBM also requires MGIDs for proper operation. When SPBM is enabled on the switch, the system reserves 519 MGIDs for SPBM operation. Therefore, the number of MGIDs on the system available for VLANs and IP multicast traffic is reduced by 519. To determine how many MGIDs are available, enter **show sys mgid-usage**.

```
ERS-8606:5# show sys mgid-usage
*****
Command Execution Time: WED FEB 02 19:42:27 2011 UTC
*****
Number of MGIDs used for VLANs : (3)
Number of MGIDs used for SPBM : (519)
Number of MGIDs used for multicast : (0)
Number of MGIDs remaining for VLANs : (1514)
Number of MGIDs remaining for multicast : (2048)
```

Before you enable SPBM on the switch, be sure that your network will not be adversely affected by this reduction in available MGIDs.

The Ethernet Routing Switch 8800/8600 supports a total of 4096 MGIDs, split between the system, VLAN, IPMC, and now SPBM. You can reserve MGIDs for IP Multicast (IPMC) traffic. You can reserve between 64 and 4084 MGIDs for IPMC. The default for IPMC is 2048. It is the responsibility of the network administrator to fully understand the network deployment strategy. Please ensure that MGIDs are planned appropriately. If assistance is required, please contact your Avaya technical representative.

For information about reserving MGIDs for IPMC, see *Avaya Ethernet Routing Switch 8800/8600 Administration* (NN46205–605).

CFM

CFM provides a mechanism to debug connectivity issues and isolate faults. This is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, the Ethernet Routing Switch 8800/8600 supports a subset of CFM functionality. CFM is based on the IEEE 802.1ag standard.

For more information about CFM, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Shortest Path Bridging MAC (SPBM)* (NN46205–525).

8800 series I/O modules

Release 7.1 introduces three new Ethernet Routing Switch 8800 interface modules. The 8800 series modules use a new enhanced network processor, the RSP 2.7.

The 8800 modules provide the same functionality as their RS module equivalents in accordance with the following table.

| RS module | New 8800 series module |
|-----------|------------------------|
| 8648GTRS | 8848GT |
| 8648GBRS | 8848GB |
| 8634XGRS | 8834XG |

For more information about 8800 series I/O modules, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Ethernet Modules* (NN46205–503).

SMLT interoperability with MSTP

This functionality allows for the connection of an SMLT/IST pair to an MSTP domain (IST switches must be root bridges). This functionality provides the ability to extend L2 VLANs from an SMLT Clustering solution towards a part of the network which is running Multiple Spanning Tree Protocol (MSTP).

There are two possible scenarios where this type of deployment might be necessary:

1. In ring topologies, it can be more cost effective to deploy MSTP to handle L2 VLANs over the fiber plant at the periphery of the network while still deploying SMLT in the Core and Data Center.
2. When migrating from a legacy Spanning Tree design towards an SMLT Clustering design with the dual requirements of being able to extend L2 VLANs between the old and new network and having a redundant design in case of link or node failures.

For more information about SMLT interoperability with MSTP, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Link Aggregation, MLT and SMLT* (NN46205–518).

4-byte AS support for BGP

In release 7.1, the Ethernet Routing Switch 8800/8600 supports both 2-byte and 4-byte AS numbers. The 4-byte AS number is a limited deployment for BGP. You can configure a BGP peer to operate in the old 2-byte AS mode or in the new 4-byte mode, not both.

For more information about 4-byte AS support for BGP, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — BGP Services* (NN46205–510).

HTTPS support for EDM

With release 7.1, the Ethernet Routing Switch 8800/8600 supports Hypertext Transfer Protocol Secure (HTTPS) connections using Enterprise Device Manager (EDM). Furthermore, after you upgrade to release 7.1, HTTPS is the default method to connect to the switch using EDM. If you require a non-secure connection (HTTP), you must disable the Web server secure-only option using CLI or ACLI. You cannot use EDM to configure HTTPS or HTTP access.

For more information about HTTPS support for EDM, see *Avaya Ethernet Routing Switch 8800/8600 User Interface Fundamentals* (NN46205-308).

RADIUS support for EDM

In Release 7.1, the Ethernet Routing Switch 8800/8600 supports RADIUS authentication of EDM connections over the web. A new option, web, has been added to the RADIUS Authentication used-by parameter. The parameter supports both HTTP or HTTPS access for EDM users.

For more information about RADIUS, see *Avaya Ethernet Routing Switch 8800/8600 Security* (NN46205–601).

EDM VRF context switching enhancement

In release 7.1, the VRF view selection no longer appears as a parameter on the embedded EDM login page. GlobalRouter is the default view. To switch to a different VRF context view, select **Configuration > VRF Context > Set VRF Context view** and use the **VRF** tab to select a different context.

The **Set VRF Context view** function is not available to RADIUS-authenticated users or to users in a service provider deployment where only a tenant VRF view is assigned. If you use a tenant VRF view, Avaya recommends that you use COM to access EDM.

For more information about VRFs, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing* (NN46205–523).

EDM navigation enhancement

In Release 7.1, EDM navigation has been enhanced. From the EDM navigation tree, you can now open all tabs with a single click rather than a double-click.

For more information about the EDM, see *Avaya Ethernet Routing Switch 8800/8600 User Interface Fundamentals* (NN46205–308).

Multiple port selection and monitoring in EDM

When you want to monitor or apply the same configuration changes to more than one port, you can use the Multiple Port selection function. If you use the embedded EDM, you can select up to a maximum of 24 ports. There is no port limit for COM users.

From the Device Physical View, you can do one of the following:

- Ctrl+click to select up to 24 specific ports.
- Click and drag to select up to 24 adjacent ports. In this case, ensure you click just outside the first port in the group and drag the mouse pointer over the group.

With both methods, selected ports appear within a yellow outline in the Device Physical View.

For more information about EDM, see *Avaya Ethernet Routing Switch 8800/8600 User Interface Fundamentals* (NN46205–308).

VRF name increase from 32 to 64 characters

Starting with Release 7.1, the Ethernet Routing Switch 8800/8600 support 64-character length names for VRFs. This feature gives you more flexibility in naming VRFs by increasing the length of this field from 32 to 64 characters.

 **Important:**

There is a backward compatibility issue with this feature because older software releases cannot recognize the 64-character VRF names. Downgrades always require previously saved configuration files (boot.cfg and config.cfg).

For more information about the VRFs, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing* (NN46205–523).

BFD enhancements

The IETF states that static routes should be taken down when BFD sessions go down. However, some equipment vendors do not comply with this standard and leave the static routes active when BFD sessions go down. Avaya allows you to choose either of the above options by implementing a static route flag. This feature enables you to control the behavior of a static route on a per VLAN/BFD session basis, which supports interoperability with all platforms and vendors.

Protocol sessions on which BFD is configured never go down regardless of the BFD session's administrative state or the administrative state of any BFD-related parameters on local or peer devices. With this enhancement, as long as there is reachability to the next hop on the BFD session, the protocols remains up and running. In previous releases, if the BFD holdoff timer was configured, the related protocol would go down when the BFD session was administratively down.

BFD is not supported with RSMLT.

For more information about the BFD, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing* (NN46205–523).

SLPP enhancements

In Release 7.1, the default SLPP protocol ID Ethertype changed from 0x8104 to 0x8102. The new Ethertype is backward compatible and supports upgrade scenarios. For example, consider two IST peers with one running Release 7.0 and the other running Release 7.1. If you set both peers to use the default SLPP Ethertype, the protocol IDs will be different but they are compatible.

For more information about the SLPP, see *Avaya Ethernet Routing Switch 8800/8600 Administration* (NN46205–605).

RSP Trace Filtering

Release 7.1 supports a new filter option for RSP trace operations.

The optional filter parameter allows you to display traces on only those RSP packets that you select. Instead of displaying all 1024 records, you can preinitialize the trace records with the values to be filtered. For example, you can filter on destination MAC addresses so that the RSP writes only MAC addresses into the trace record.

For more information about RSP Trace Filtering, see *Ethernet Routing Switch 8800/8600 Troubleshooting* (NN46205-703).

show command timestamp

Release 7.1 supports an enhancement that displays a timestamp on the output of show commands. With this feature, you can compare the output of commands executed at different times, which can help in debugging. The timestamp feature is enabled by default.

CP limit statistics

CP Limit protects the CPU from being flooded by traffic from a single, unstable port. You configure CP Limit by specifying thresholds on specified port within the chassis. If an unstable port reaches this threshold, CP Limit logs the current port statistics and then shuts down the port.

Release 7.1 supports the collection of statistics for CP Limit. The CP Limit statistics feature captures traffic details such as the type of traffic and their queue priority. This information helps in the debugging of network issues.

For more information about CP limit statistics, see *Ethernet Routing Switch 8800/8600 Administration* (NN46205-605).

CLI wrapper for shell debug commands

The shell debug wrapper commands group multiple shell commands under one command. With this feature, you can enter one command to display the output of all the related shell

commands in that group to facilitate debugging. The show debug wrapper commands are separated into the following groups:

- generic (general shell and CPP commands)
- platform
- bridging
- routing
- multicast
- spbm (Shortest Path Bridging commands)

For more information about debug commands, see *Ethernet Routing Switch 8800/8600 Troubleshooting* (NN46205-703).

Flash file system enhancements

Release 7.1 supports the `dos-stop` command to ensure that the external compact flash card is synchronized before it is removed to avoid removed flash card issues.

Release 7.1 also supports the `shutdown` command to ensure that the SF/CPU card is shutdown before it is removed or the system is powered off.

For more information about Flash file system enhancements, see *Avaya Ethernet Routing Switch 8800/8600 — Administration* (NN46205-605).

Chapter 3: Important notices

This section describes the supported and unsupported hardware and software features in the Avaya Ethernet Routing Switch 8800/8600 Software Release 7.1, and provides important information for this release.

Navigation

- [Supported hardware and software compatibility](#) on page 18
- [Unsupported hardware for Release 7.1](#) on page 21
- [Supported software and hardware scaling capabilities](#) on page 22
- [Software licensing](#) on page 26
- [File names for this release](#) on page 27
- [Important information and restrictions](#) on page 29
- [High Availability mode considerations](#) on page 47
- [Ongoing considerations](#) on page 48
- [Module and chassis compatibility and performance considerations](#) on page 48
- [Switch clustering topologies and interoperability with other products](#) on page 50
- [SF/CPU protection and loop prevention compatibility](#) on page 50
- [Switch behavior during boot cycle and redundant configuration files](#) on page 50
- [MPLS considerations](#) on page 54
- [SNMP considerations](#) on page 54
- [DVMRP considerations](#) on page 55
- [SMLT considerations](#) on page 55
- [RSMLT considerations](#) on page 55
- [IST considerations](#) on page 56
- [60 day trial license](#) on page 56
- [Advanced filter guidelines](#) on page 57
- [MTBF for 1 Gig SFPs](#) on page 57

- [Supported standards, RFCs, and MIBs](#) on page 58
- [Supported traps and notifications](#) on page 58

Supported hardware and software compatibility

The following table describes your hardware and the minimum Ethernet Routing Switch 8800/8600 software version required to support the hardware.

Table 1: Chassis, power supply, and SF/CPU compatibility

| Item | | Minimum software version | Part number |
|---|--|--------------------------|--------------------------------|
| Chassis | | | |
| 8010co | 10-slot | 3.1.2 | DS1402004-E5 DS1402004-E5GS |
| 8010 | 10-slot | 3.0.0 | DS1402001-E5 DS1402001-E5GS |
| 8006 | 6-slot | 3.0.0 | DS1402002-E5 DS1402002-E5GS |
| 8003-R | 3-slot | 7.0.0.0 | DS1402011-E5 |
| Switching fabric/CPU | | | |
| 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 |
| 8895 SF/CPU | Switching fabric | 7.0 | DS1404120-E5 |
| Power supplies | | | |
| 8004AC | 850 W AC | 3.1.2 | DS1405x08 |
| 8004DC | 850 W DC | 3.1.2 | DS1405007 |

| Item | | Minimum software version | Part number |
|----------------|----------------------|--------------------------|--------------|
| Chassis | | | |
| 8005AC | 1462 W AC | 4.0.0 | DS1405012 |
| 8005DI AC | 1462 W Dual input AC | 5.0 | DS1405018-E6 |
| 8005DI DC | 1462 W Dual input DC | 5.1 | DS1405017-E5 |
| 8005DC | 1462 W DC | 4.0.x | DS1405011 |

Table 2: Module and component compatibility

| Modules and components | | Minimum software version | Part number |
|----------------------------|--|--------------------------|--------------|
| Ethernet R modules | | | |
| 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 |
| Ethernet RS modules | | | |
| 8848GB | 48 100/1000 Mbps SFP ports | 7.1.0.0 | DS1404122-E6 |
| 8848GT | 48-port 10/100/1000 Mbps copper ports | 7.1.0.0 | DS1404123-E6 |
| 8834XG | 24 100/1000 Mbps SFP ports 2 XFP ports 8 10/100/1000 Mbps copper ports | 7.1.0.0 | DS1404124-E6 |
| 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 2 XFP ports 8 10/100/1000 Mbps copper ports | 5.0.0 | DS1404109-E6 |
| 8648GBRS | 48 100/1000 Mbps SFP ports | 5.0.0 | DS1404102-E6 |

| Modules and components | | Minimum software version | Part number |
|---|--|--------------------------|------------------------------|
| Small form factor pluggable transceivers | | | |
| 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 |
| 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, up to 10 km | 4.1.0 | AA1419069-E6 |

| Modules and components | | Minimum software version | Part number |
|---|--|--------------------------|--------------|
| 1000BASE-BX bidirectional SFP | 1490 nm, single fiber LC, up to 10 km | 4.1.0 | AA1419070-E6 |
| 1000BASE-BX bidirectional SFP | 1310 nm, single fiber LC , up to 40 km | 7.0 | AA1419076-E6 |
| 1000BASE-BX bidirectional SFP | 1490 nm, single fiber LC, up to 40 km | 7.0 | AA1419077-E6 |
| 1000BASE-EX | 1550 nm, up to 120 km | 5.0 | AA1419071-E6 |
| 10 Gigabit Ethernet Small form factor pluggable transceivers | | | |
| 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 |
| 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 |

Unsupported hardware for Release 7.1

Release 7.1 does not support any classic modules, including the following:

- 8608GBE module
- 8608GBM module
- 8608GTE module
- 8608GTM module
- 8608SXE module
- 8616GTE module
- 8616SXE module
- 8624FXE module
- 8632TXE module
- 8632TXM module

- 8648TXE module
- 8648TXM module
- 8672ATME module
- 8672ATMM module
- 8683POSM module
- 8690 SF/CPU module
- 8691 SF/CPU module
- Web Switching Module (WSM)
- 8660 Service Delivery Module (SDM)
- 8661 SSL Acceleration Module (SAM)
- Media Dependent Adapters for the 8672ATME and 8672ATMM Modules
- Breaker Interface Panel
- 8001AC power supply
- 8002DC power supply
- 8003AC power supply

Release 7.1 supports the 8692 SF/CPU only if it is equipped with SuperMezz. The 8692 SF/CPU without SuperMezz is not supported with Release 7.1.

In addition, M mode is no longer supported in Release 7.1. The software runs in R mode by default.



Important:

In release 7.1, the 8003 chassis is no longer supported. It is replaced by the 8003-R chassis.

Supported software and hardware scaling capabilities

The following table lists supported software and hardware scaling capabilities in Ethernet Routing Switch 8800/8600 Software Release 7.1. The information in this table supersedes information contained in *Avaya Ethernet Routing Switch 8800/8600 Planning and Engineering — Network Design, NN46205-200*, or any other document in the suite.

The capabilities described in this table were tested as individual protocols, not mixtures of protocols.

Avaya 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 Avaya 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).

Table 3: Supported scaling capabilities

| | Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU |
|---|--|
| <i>Shortest Path Bridging MAC (SPBM)</i> | |
| Multicast scaling | 2000 |
| MAC scaling routed | 6000 |
| ARP scaling routed | 6000 |
| VLAN scaling routed | 100 |
| VRF instances | 500 |
| NRE routes | 8000 |
| VRF routes total | 8000 |
| L2 VPNs | 1000 |
| Number of MACs onVPNs | 30 000 |
| <i>Layer 2</i> | |
| MAC address table entries | 64 000 32 000 when SMLT is used |
| VLANs (port- protocol-, and IEEE 802.1Q-based) | 4000 |
| IP subnet-based VLANs | 800 |
| Ports per Link Aggregation Group (LAG, MLT) | 8 |
| Aggregation groups 802.3ad aggregation groups Multi Link Trunking (MLT) group | 128 |
| SMLT links | 128 |
| SLT (single link SMLT) | 382 |
| VLANs on SMLT/IST link | With Max VLAN feature enabled: 2000 |
| RSMLT per VLAN | 32 SMLT links with RSMLT-enabled VLANs |
| RSTP/MSTP (number of ports) | 384, with 224 active. Configure the remaining interfaces with Edge mode |
| MSTP instances | 32 |
| <i>Advanced Filters</i> | |
| ACLs for each system | 4000 |

Important notices

| | Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU |
|--|---|
| ACEs for each system | 10 000 |
| ACEs for each ACL | 1000 |
| ACEs for each port | 2000: 500 inPort 500 inVLAN 500 outPort 500 outVLAN |
| <i>IP, IP VPN/MPLS, IP VPN Lite, VRF Lite</i> | |
| IP interfaces (VLAN- and router-based) | 1972 |
| VRF instances | 255 |
| ECMP routes | 5000 |
| VRRP interfaces | 255 |
| IP forwarding table (Hardware) | 250 000 |
| BGP/mBGP peers | 250 |
| iBGP instances | on GRT |
| eBGP instances | on 256 VRFs (including GRT) |
| BGP forwarding routes BGP routing information base (RIB) BGP forwarding information base (FIB) | BGP FIB 250 000 BGP RIB 500 000 |
| IP VPN routes (total routes for each system) | 180 000 |
| IP VPN VRF instances | 255 |
| Static ARP entries | 2048 per VRF 10 000 per system |
| Dynamic ARP entries | 32 000 |
| DHCP relay instances (total for all VRFs) | 512 |
| Static route entries | 2000 per VRF 10 000 per system |
| OSPF instances for each switch | on 64 VRFs (including GRT) |
| OSPF areas for each switch | 5 per VRF 24 per system |
| OSPF adjacencies for each switch | 80 200 per system |
| OSPF routes | 20 000 per VRF 50 000 per system |
| OSPF interfaces | 500 500 per system |
| OSPF LSA packet maximum size | 6000 bytes |
| RIP instances | on 64 VRFs (including GRT) |
| RIP interfaces | 200 |
| RIP routes | 2500 per VRF 10 000 per system |
| <i>Multiprotocol Label Switching (MPLS)</i> | |

| | Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU |
|--|---|
| MPLS LDP sessions | 200 |
| MPLS LDP LSPs | 16 000 |
| MPLS RSVP static LSPs | 200 |
| Tunnels | 2500 |
| <i>IP Multicast</i> | |
| DVMRP passive interfaces | 1200 |
| DVMRP active interfaces/neighbors | 80 |
| DVMRP routes | 2500 |
| PIM instances | on 64 VRFs (including GRT) |
| PIM active interfaces | 200 (200 for all VRFs) |
| PIM passive interfaces | 1972 (2000 for all VRFs) |
| PIM neighbors | 80 (200 for all VRFs) |
| MSDP peers | 20 |
| MSDP maximum SA messages | 6144 |
| Multicast streams: with SMLT/ without SMLT | 2000/4000 |
| Multicast streams per port | 1000 |
| IGMP reports/sec | 250 |
| <i>IPv6</i> | |
| IPv6 interfaces | 250 |
| IPv6 tunnels | 350 |
| IPv6 static routes | 2000 |
| OSPFv3 areas | 5 |
| OSPFv3 adjacencies | 80 |
| OSPFv3 routes | 5000 |
| <i>Operations, Administration, and Maintenance</i> | |
| IPFIX | 384 000 flows per chassis |
| RMON alarms with 4000K memory | 2630 |
| RMON events with 250K memory | 324 |
| RMON events with 4000K memory | 5206 |
| RMON Ethernet statistics with 250K memory | 230 |

| | Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU |
|--|--|
| RMON Ethernet statistics with 4000K memory | 4590 |

Software licensing

The following table describes the license required to use specific features. The Premier License enables all licensed features on the Ethernet Routing Switch 8800/8600.

Table 4: License and features

| Base License | Advanced License | Premier License |
|---|---|--|
| <ul style="list-style-type: none"> • IP Multinetting • IP Source Guard • DHCP Snooping • Dynamic ARP Inspection • BPDU Filtering • IGMP Querier for L2 • PIM-SSM for SMLT • Multicast VLAN Registration (MVR) | <ul style="list-style-type: none"> • all Base License features • Border Gateway Protocol version 4 (BGPv4) for more than 10 Peers • Bidirectional Forwarding Detection (BFD) • Multicast Source Discovery Protocol (MSDP) • Packet Capture function (PCAP) • IPv6 Features: <ul style="list-style-type: none"> - IP Routing - IPv6 over SMLT and RSMLT - DHCPv6 Relay - VRRPv3 - BGP+ - RADIUSv6 | <ul style="list-style-type: none"> • all Base License and Advanced License features • Virtual Routing and Forwarding Lite (VRF Lite) • Multi-Protocol Border Gateway Protocol (MP-BGP) • IP-Virtual Private Network, Multi-Protocol Label Switching (RFC2547) (IP-VPN MPLS RFC2547) • IP-Virtual Private Network-Lite (IP-VPN-Lite – IP in IP) • Multicast virtualization for VRF-Lite (IGMP and PIM-SM/SSM) • Shortest Path Bridging (SPB) Features: <ul style="list-style-type: none"> - SPB L2 VSNs (VLAN Extensions) - SPB Native IP shortcuts (VRF0 shortcuts) - SPB L3 VSNs (VRF Extensions) - IP VPN Lite over SPB IP shortcuts |

| Base License | Advanced License | Premier License |
|--------------|------------------|--|
| | | <ul style="list-style-type: none"> - Inter-ISID Routing - IEEE 802.1ag Connectivity Fault Management |

All IPv6 features require the Advanced License.

Ethernet Routing Switch 8800/8600 Release 7.1 includes a Premier trial license that is valid for 60 days from the date of install. After 60 days, the license expires and configured licensed features are no longer functional after the switch is restarted or rebooted. If you want these configured features to continue to function properly, you must install a valid license.

For more information about using licenses, see *Avaya Ethernet Routing Switch 8800/8600 Administration* (NN46205-605).

File names for this release

This section describes the Ethernet Routing Switch 8800/8600 Software Release 7.1 software files.

Before you upgrade, Avaya recommends that you verify the MD5 signature for each new file to be used. For upgrade procedures, see *Avaya Ethernet Routing Switch 8800/8600 Upgrades — Software Release 7.1, NN46205-400*.

Table 5: Release 7.1 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_7100.tar.gz | 64,392,948 |
| Copyright file | Ethernet Routing Switch 8600/8800 Master Copyright file | ERS8k.7.1.0.0_Copyright.docx | 56,205 |
| Ethernet Routing Switch images | | | |
| Boot monitor image for 8692 SF/CPU | 8692 CPU and switch fabric firmware | p80b7100.img | 1,186,330 |
| Boot monitor image for 8895 SF/CPU | 8895 CPU and switch fabric firmware | p80be7100.img | 1,254,143 |

| Module or file type | Description | File name | Size in bytes |
|---|--|--|---------------|
| Run-time image for 8692 SF/CPU | Run-time image for 8692 SF/CPU | p80a7100.img | 15,366,695 |
| Run-time image for 8895 SF/CPU | Run-time image for 8895 SF/CPU | p80ae7100.img | 14,515,839 |
| Run-time image for R modules | Image for R modules | p80j7100.dld | 1,743,992 |
| Run-time image for RS modules | Run-time image for RS modules | p80k7100.dld | 1,807,392 |
| Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz) | Image for the SuperMezz card | p80m7100.img | 15,468,016 |
| 3DES for 8692 SF/CPU | Encryption module for privacy protocol with Secure Shell (SSH) | p80c7100.des | 56,124 |
| 3DES for 8895 SF/CPU | Encryption module for privacy protocol with Secure Shell (SSH) | p80ce7100.des | 51,972 |
| AES for 8692 SF/CPU | Encryption module for privacy protocol with Secure Shell (SSH) | p80c7100.aes (this image includes the 3DES image) | 27,436 |
| AES for 8895 SF/CPU | Encryption module for privacy protocol with Secure Shell (SSH) | p80ce7100.aes (this image includes the 3DES image) | 25,156 |
| MIB | MIB files | p80a7100.mib | 5,229,393 |
| MIB (zip file) | Zip file containing MIBs | p80a7100.mib.zip | 819,015 |
| MD5 checksum file | md5 checksums of all Release 7.1 software files | p80a7100.md5 | 1,224 |
| Firmware images | | | |
| FOQ for R modules | Feedback output queueing FPGA firmware | foq267.xsvf | 5,320,469 |

| Module or file type | Description | File name | Size in bytes |
|--|---|------------------------|---------------|
| BMC for R modules | BAP memory controller FPGA firmware | bmc776.xsvf | 2,640,266 |
| DPC for R modules | Dual port controller FPGA firmware | dpc194.xsvf | 2,642,001 |
| PIM8630GBR | Programmable I/O module FPGA firmware; for the 8630GBR only | PI_769.xsvf | 2,284,578 |
| Firmware for RS modules | Contains FOQ, BMC, DPC, mirroring, and loopback images | rs_dpm_fpga.bin | 4,538,368 |
| PIM images for RS modules | PIM FPGA firmware required for 8612XLRS module only | pim8612XLRS.bin | 60,183 |
| | PIM FPGA firmware required for 8634XGRS module only | pim8634XGRS.bin | 78,173 |
| | PIM FPGA firmware required for 8648GBRS module only | pim8648GBRS.bin | 79,891 |
| | PIM FPGA firmware required for 8648GTRS module only | pim8648GTRS.bin | 54,441 |
| Trace files | | | |
| MPLS trace file | Trace file for MPLS. This is autogenerated and appears on the PCMCIA after upgrade. | nbpdtrc.lo0 | variable |
| EDM Help files | | | |
| EDM help files | Help files for EDM GUI | ERS8k.7.1.0.0_Help.zip | 2,973,741 |
| ERS 8000/8600 EDM plug-in for COM | | | |
| EDM plug-in for COM | EDM plug-in for COM | ers8000v7.1.0.0.war | 7,463,823 |

Important information and restrictions

This section contains important information and restrictions that you should consider before you upgrade to Release 7.1.

Fixes from previous releases

The Ethernet Routing Switch 8800/8600 Software Release 7.1 incorporates all fixes from prior releases up to and including release 5.1.4.0.

Important information and restrictions navigation

- [SuperMezz, SF/CPU memory, and upgrades](#) on page 30
- [Compact flash card display on 8895 SF/CPU](#) on page 31
- [Proper care of external compact flash and PCMCIA cards](#) on page 31
- [EDM considerations](#) on page 32
- [Installing EDM help files](#) on page 35
- [I/O module considerations](#) on page 36
- [MLT/LAG considerations](#) on page 36
- [Console connection considerations](#) on page 36
- [DHCP snooping considerations](#) on page 36
- [MLT/LAG considerations](#) on page 36
- [Supported upgrade paths](#) on page 37
- [General upgrade considerations](#) on page 37
- [Upgrade considerations for Release 7.1](#) on page 37
- [Upgrading from 8692 SF/CPU with SuperMezz to 8895 CPU](#) on page 38
- [Upgrade considerations: DOSFS with upgrades from pre-Release 5.0](#) on page 39
- [Upgrade considerations: Power Management](#) on page 40
- [Upgrade considerations: IST](#) on page 43
- [Pre-release 5.1 upgrades considerations: specifying license file location](#) on page 43
- [Considerations for upgrades from 5.0-based code releases](#) on page 44

SuperMezz, SF/CPU memory, and upgrades

To support Release 7.1, the 8692 SF/CPU must be equipped with SuperMezz. 8692 SF/CPU without SuperMezz is not supported with Release 7.1. If the Release 7.1 software is booted with a non-SuperMezz 8692 SF/CPU, the line cards do not come online.

For Release 7.1, Avaya recommends that the PCMCIA card for the 8692 SF/CPU with SuperMezz be at least 256 MB. 256 MB is the current size of the shipping PCMCIA card. The 8692 SF/CPU with SuperMezz does not support PCMCIA cards larger than 256 MB.

The 8895 SF/CPU comes with a 2 GB compact flash card.

Compact flash card display on 8895 SF/CPU

The 8692 SF/CPU with SuperMezz displays the external PCMCIA card as `/pcmcia` . The 8895 SF/CPU has an external compact flash card installed rather than a PCMCIA card, and also displays this flash card as `/pcmcia` .

The internal flash memory (64 MB) is displayed as `/flash` for both the 8692 SF/CPU with SuperMezz and the 8895 SF/CPU.

Proper care of external compact flash and PCMCIA cards

To guarantee the external compact flash card or the PCMCIA card is in a consistent state before you remove it, use one of the following commands.

- `pcmcia-stop` (on 8692 SF/CPU)
- `dos-stop /pcmcia` (on 8895 SF/CPU)

Do not remove the external memory card without first entering one of the preceding commands.

Be sure to back up all configurations, as all files can be lost if the card becomes corrupted.

To check and optionally repair a file system, you can use the `dos-chkdsk <device> repair` command.

If the file system cannot be repaired, you can attempt to reformat the device using the `dos-format <device>` command. Otherwise, you may need to replace the card.

Both of the above commands delete all information on the memory, so be sure to backup all information before using either of the commands.

The above commands are available in the CLI, ACLI, or the boot monitor.

Pasting configurations into the configuration file

If you use the console, Telnet, or SSH to paste configurations into the switch configuration file, use the following guidelines:

- Use an ASCII-only editor and do not include any additional (hidden) characters.
- Make sure that the order of the commands is correct.

EDM considerations

In the EDM Physical Device view, EDM does not display the name of the 8692 SF/CPU cards. This issue does not affect 8895 SF/CPU cards.

In EDM, if you create a BGP Peer (under **Configuration > IP > BGP > Peers > Insert**), the AdvertisementInterval value defaults to 30. This value should default to 5, which is the default route advertisement interval value for configuration using the CLI or ACLI.

The following sections list other EDM considerations.

Supported browsers

For Enterprise Device Manager (EDM) to display and function correctly, use one of the following Web browsers:

- Mozilla Firefox, version 3.0+
- Microsoft Internet Explorer, version 7.0

If you connect to EDM using an unsupported browser, the switch displays an error message.

On-box and off-box EDM

EDM is a Web-based graphical user interface (GUI) for element management and configuration of the Ethernet Routing Switch 8800/8600. EDM is an embedded application on the Ethernet Routing Switch, and the EDM Web server is the switch itself.

EDM for the Ethernet Routing Switch 8800/8600 is also supported as a plug-in with the Configuration and Orchestration Manager (COM). Access to COM is also through a browser.

To distinguish between the embedded EDM and the EDM plug-in for COM, the following terminology is used in the Ethernet Routing Switch 8800/8600 documentation:

- on-box EDM: EDM software that is embedded with the switch code
- off-box EDM: EDM plug-in that is available with the COM software

Note:

If you launch on-box EDM using Internet Explorer and then graph a port, you cannot change the default 5s polling interval from the drop down box. As a workaround, you can launch on-box EDM using Firefox, or use the off-box EDM plug-in.

Saving runtime configurations in EDM

In EDM, the option for saving runtime configuration changes is not easily seen. To save current changes, go to **Configuration > Edit > Chassis** and under ActionGroup1, click on **SaveRuntimeConfig** and click **Apply**. (Q02114591)

Unlike Java Device Manager, when you exit EDM, there is no pop-up dialog box prompting you to save the configuration.

EDM table display

Avaya does not recommend using EDM (on-box or off-box through COM plug-in) to display routing tables with 3000 or more entries as doing so can take a long period of time (many minutes) to formulate the display. The EDM application can become unusable until the whole table is displayed. This issue is present with all large route tables, but is more apparent with BGP route tables. Avaya recommends that you use either the CLI or ACLI to display these type of tables. Be aware that this display scenario does not affect traffic on the switch.

This same recommendation previously applied to Java Device Manager operations. (Q02123849)

EDM functionality differences from Java Device Manager

In some cases, EDM functionality differs from that previously offered in Java Device Manager (JDM), including the following:

- **Single username and password combination for each VRF**

With EDM, you can configure only one username and password combination for each VRF.

- **Managing VRF users with COM**

With COM, Avaya recommends that the administrator of the COM system assign appropriate device credentials along with proper VRF mapping to COM users.

- If a COM user needs to be restricted to a particular VRF, in the device credentials, map the credentials for the COM user to that VRF.
- If a COM user needs GlobalRouter access, in the device credentials, map the credentials for the COM user to the GlobalRouter. GlobalRouter access allows the COM user access to any and all VRFs.

Upon launching the EDM plugin, users with restricted VRF can see the device view for that particular VRF only. Users with the GlobalRouter VRF associated have the ability to switch the VRF context to another VRF as needed.

 **Important:**

In COM, the VRF Manager allows you to further restrict access to a device to a particular VRF. When you launch the EDM plugin, the displayed VRF is the one specified by the VRF Manager (assuming the appropriate user credentials are also configured). However, in the case where your user credentials are mapped to the GlobalRouter, and the VRF Manager maps the device to a specific VRF, the EDM plugin launches the specified non-GlobalRouter VRF rather than the GlobalRouter VRF. Furthermore, in this scenario, you cannot switch the VRF context to another VRF using the EDM plugin.

As a result, to switch the VRF context, Avaya recommends that you not use the VRF Manager to map the VRF to a non-GlobalRouter VRF. Instead, map the VRF to the GlobalRouter in the VRF Manager, and use the Set VRF menu option from within the EDM off-box plugin (described above) to switch the device context to a different VRF.

If a COM user finds an unexpected behavior with an incorrect default VRF context being launched for the EDM plugin inside COM, do the following:

- Check the credentials in COM for that device. To access credentials, in the COM left panel, expand **Admin** and click **Device Credentials**. Verify that the COM user is assigned the correct VRF (to allow the user to switch between multiple VRF contexts, they must be assigned to VRF 0 or GlobalRouter).
- If the credentials are correct, check the VRF manager in COM. In the COM left panel, expand **Managers** and click the **Virtual Routing Manager** icon. Make sure that the device has the correct VRF associated with it (VRF 0 or GlobalRouter to allow the user to switch between multiple VRF contexts). If a device is assigned a specific VRF in the VRF Manager, all functions within COM (including EDM) use that VRF context by default.

Also be aware of the following:

- In order to modify the VRF context using the VRF Manager, the user needs GlobalRouter credentials for a device in the device credentials page.
- The VRF Manager is available in COM only if the full COM application license is purchased.
- The VRF Manager must be assigned to a particular user by the COM administrator using the Manager assignment function under the Admin/Access Control menu in the COM left navigation pane. This option exists in order to allow role-based access control for users to whom the administrator wishes to limit privileges when there are many users of the system.

• **CLI window launch**

The on-box EDM GUI is a browser-based solution that can run from any supported platform (Windows or Linux) and it does not offer the capability to launch a Windows-based command prompt window as was available in JDM. In the COM with off-box Ethernet Routing Switch 8800/8600 EDM plug-in, the CLI manager exists to launch CLI

windows as needed. You can also connect to a switch using your own local command prompt.

- **Supported COM release**

For Release 7.1, Avaya recommends using COM 2.2 or higher.

Using the EDM plug-in with COM

The Configuration and Orchestration Manager (COM) is an Avaya off-box network management tool that supports an EDM plug-in for the Ethernet Routing Switch 8800/8600. The EDM plug-in allows you to perform EDM functions within the off-box COM tool. For information about installing the EDM plug-in for COM, see *Avaya Configuration and Orchestration Manager Using the Product Interfaces* (NN47226-100).

You can obtain the EDM plug-in software from the Avaya support site at <http://support.avaya.com>.

Installing EDM help files

While the EDM GUI is bundled with the Release 7.1 software, the associated EDM help files are not included. To access the help files from the EDM GUI, you must install the EDM help files on either a TFTP or FTP server in your network.

Use the following procedure to install the EDM help files on a TFTP or FTP server.



Important:

Do not install the EDM help files within the `/pcmcia` or `/flash` file systems, as the help files consume too much space.

Procedure steps

1. Retrieve the EDM help zip file from avaya.com or from the software CD.
2. On a TFTP or FTP server that is reachable from your 8800/8600 switch, create a directory named: `ERS8000_71_Help`.

If you are using FTP for this installation, be sure that the 8800/8600 switch is configured with the appropriate host name and password using the `config bootconfig host user` and `config bootconfig host password` commands (or, using the ACLI, `boot config host user` and `boot config host password`).

If a host password is configured, the 8800/8600 switch uses FTP to transfer data from the switch to the server. If no host password is configured, the switch uses TFTP for the data transfer. To clear the host password, specify a blank value using the host password command: `config bootconfig host password ""` (CLI) **OR** `boot config host password ""` (ACLI)

3. Unzip the EDM help zip file in the new FTP or TFTP server directory.
4. Using EDM on the 8800/8600 switch, open the following folders: **Configuration, Security, Control Path.**
5. Double-click **General.**
6. Click the **Web** tab.
7. In the **HelpTftp/Ftp_SourceDir** field, enter the FTP or TFTP server IP and the path of the online directory where the files are unzipped, in the following format: `<TFTP/FTP-server-IP-address>:ERS8000_70_Help.`
8. To test that the help is working properly, select any tab (for example, **Edit > Chassis**) and click the **Help** button.

The appropriate EDM help page appears.

I/O module considerations

The 8648GTR module does not support a packet size larger than 9188 bytes at 100 Mbps. At 1000 Mbps, frames larger than 9188 bytes (up to 9600 bytes) are supported.

MLT/LAG considerations

To maintain MLT and LAG stability during failover, Avaya recommends the use of CANA: you must configure the advertised speed to be the same for all MLT/LACP links. For 10/100/1000 Mbps ports, ensure that CANA uses only one specific setting, for example, 1000-full or 100-full. Otherwise, a remote device could restart Auto-Negotiation and the link could use a different capability. In the case of LACP LAGs, ports of different speeds cannot join the same LAG.

It is important that each port uses only one speed and duplex mode. The use of CANA forces this setting. 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/LAG.

Console connection considerations

If you change the management IP setting using EDM or an SNMP device, the active console session is terminated. In this case, you must reopen the console session.

DHCP snooping considerations

On any switch configured with both DHCP Relay and DHCP snooping enabled, you must ensure that the routing interfaces where the DHCP offer is received are configured as DHCP

snooping trusted ports. This applies to any and all return paths; that is, primary and backup routing interfaces.

Supported upgrade paths

The Ethernet Routing Switch 8800/8600 Software Release 7.1 supports direct upgrades from the following earlier releases:

- 4.1.8
- 5.0.1
- 5.1.2
- 5.1.3
- 5.1.4
- 7.0

If you want to upgrade to release 7.1 from any other release, first upgrade to one of the above releases and then upgrade to 7.1.

General upgrade considerations

The configuration file generated with Ethernet Routing Switch 8800/8600 Software Release 7.1 contains options that are not backward-compatible with any previous Ethernet Routing Switch 8800/8600 Software Releases.

Loading a Release 7.1 configuration file on a pre-7.1 runtime image can generate errors and cause the image to stop loading the configuration file. Under these conditions, the system will load with a default configuration.

If 8800/8600 switches running pre-7.0 code are connected to rebranded 8800 7.0 switches, the pre-7.0 switches cannot identify the chassis type and remote port from Topology Discovery Packets from the rebranded 8800 switches. As a result, in the pre-7.0 switches, the command **show sys topology** displays `unknown error: 192` in the ChassisType and Rem Port fields for the 8800 switches.

Downgrades always require previously saved configuration files (boot.cfg and config.cfg) and may require the removal of R, RS, and 8800 series modules prior to downgrade.

Upgrade considerations for Release 7.1

Before you upgrade, read *Avaya Ethernet Routing Switch 8800/8600 Upgrades — Software Release 7.1, NN46205-400* and follow the outlined procedures.

If you are upgrading from a release prior to 5.0, you must 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](#) on page 39.

You must take into consideration Power Management for this release; for more information, see [Upgrade considerations: Power Management](#) on page 40.

Upgrading from 8692 SF/CPU with SuperMezz to 8895 CPU

Use the following steps to upgrade from 8692 SF/CPU with SuperMezz to 8895 CPUs.

Prerequisites

- You must be local to the switch with a console connection.
- Upgrade the Ethernet Routing Switch 8800/8600 to 7.1 code with the 8692 SF/CPU with SuperMezz as master and slave.
- Download the p80ae7100.img and p80be7100.img software images, as well as the dld files (p80j7100.dld, p80k7100.dld) to the master 8692 SF/CPU.

Procedure steps

1. Disable the slot for the slave SF/CPU. For example:

```
ERS-8010:5# config slot x state dis (where slot x is the slot of the slave 8692 SF/CPU).
```

2. Remove the slave 8692 SF/CPU with SuperMezz.
3. Insert the 8895 SF/CPU into the chassis, and immediately after inserting the 8895 SF/CPU, stop the boot process at the boot monitor when prompted.
4. Copy the running configuration file (config.cfg), boot configuration file (boot.cfg), images and dld files (p80ae7100.img, p80be7100.img, p80j7100.dld, p80k7100.dld) from the current master 8692 SF/CPU to the 8895 SF/CPU using the internal IP for the copy command: 127.0.0.X, where X is the slot number of the 8692 SF/CPU. For example:

```
ERS-8010:5# copy 127.0.0.X:/flash/<name of the file> /flash/
```

5. Edit the primary image file name in the boot.cfg to load the 8895 image. For example:

```
monitor:5# choice primary image-file p80ae7100.img monitor:5#  
save
```

6. Boot the 8895 SF/CPU with the correct image and wait for the login screen. For example:


```
monitor:5# boot /flash/ p80be7100.img
```
7. Perform a failover from the master 8692 SF/CPU using the following command:


```
config sys set action cpuswitchover
```
8. After the 8895 SF/CPU becomes the master, remove the slave 8692 SF/CPU with SuperMezz.
9. Insert another 8895 SF/CPU into the chassis, and immediately after inserting the 8895 SF/CPU, stop the boot process at the boot monitor when prompted.
10. Copy the running configuration file (config.cfg), boot configuration file (boot.cfg), images and dld files (p80ae7100.img, p80be7100.img, p80j7100.dld, p80k7100.dld) from the current master 8895 SF/CPU to the new 8895 SF/CPU using the internal IP for the copy command: 127.0.0.X, where X is the slot number of the master 8895 SF/CPU. For example:


```
ERS-8010:5# copy 127.0.0.X:/flash/<name of the file> /flash/.
```
11. Boot the 8895 SF/CPU with the correct images and wait for the login screen.


```
monitor:5# boot /flash/ p80be7100.img
```

Upgrade considerations: DOSFS with upgrades from pre-Release 5.0

Release 5.0 introduced 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 7.1, 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. Avaya recommends backing up the files from /flash, and executing dos-format /flash to bring the file system on the /flash device to the latest ERS 8800/8600 baseline.
```

If you receive this message, Avaya 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. Avaya recommends that you back up all files from the DOSFS device, reformat the device, and replace all files.

Be sure to back up hidden files as well. For information about hidden files, see *Avaya Ethernet Routing Switch 8800/8600 Upgrades — Software Release 7.1* (NN46205-400).

Upgrade considerations: Power Management

The Power Management feature available with Release 7.1 may require you to take special steps before you upgrade.

When you upgrade to Release 7.1, 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 8800/8600 system. To determine the number of power supplies required for your switch configuration, use the *Avaya ERS 8800/8600 Power Supply Calculator, NN48500-519*. This is available on the Avaya support Web site at www.avaya.com/support.

For Power Management configuration and conceptual information, see *Avaya Ethernet Routing Switch 8800/8600 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.

No I/O modules are 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.

The amount of system power is calculated based on the number, type, and input source voltage of the power supplies in the chassis. This system power calculation is equal to the DC wattage output (which can differ depending on AC input voltage) minus 90 W required for the fans. 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 8004 series power supplies, the system power output calculation is the same (690 W), regardless of source input AC voltage. However, the actual power supply wattage output will vary depending upon the input source voltage. The system power output calculation is always based on low-voltage input. Therefore in systems using 8004 series power supplies that are running at high voltage input (220 V), the system output power calculation will actually be lower (displaying 690 W) than what the system is capable of.

By default, switch fabrics are allotted highest priority and always power up. I/O modules power up if there is sufficient power remaining to do so. If there is insufficient power to bring all I/O modules online, they are powered up based on slot priority. By default, I/O modules are powered up starting at slot 1 until there is insufficient power to bring the next module online.

You have the ability within a working system to reconfigure slot priority to your own requirements. Avaya does not recommend changing the priority for the switch fabric slots.

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 after the additional power supply is added.

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. Correct this situation as soon as possible.

Disabling power and cooling management

You can 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 8800/8600 system. To determine the number of power supplies required for your switch configuration, use the *Power Supply Calculator for Avaya ERS 8800/8600, NN48500-519*. This is available on the Avaya support Web site at www.avaya.com/support.

Important:

Avaya 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.

By default, RS modules do not come up when the high-speed cooling module is not installed.

Important:

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

Use the following procedure in order to override the fan check for the high-speed cooling modules.

1. Save the pre-7.1 or current 7.1 configuration file.

```
save <file-name>.cfg
```

2. Edit the configuration file offline using an editor like VI or EMACS. You can either:
 - Use the CLI to edit the file on the switch (the switch has a built-in VI-like editor). Use the `edit <file-name>.cfg` command.
 - Save the file as an ASCII file and transfer to another device for editing with a text editor like Notepad.
 - Transfer the file to a device and edit with VI or an EMACS-like editor, or using a text editing application such as MS Word. The configuration file is plain text only.
3. In the configuration file, add the following lines to the end of the flags section:


```
#!power power-check-enable false
#!power fan-check-enable false
```

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.
5. Reboot the switch or source the configuration file.

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.

Pre-release 5.1 upgrades considerations: specifying license file location

If you upgrade to release 7.1 from a release prior to 5.1, you must specify the location of your license file in the boot configuration file. If you do not specify the location of your license file, you can encounter issues with your licensed features.

Procedure steps

To specify the license file location, enter the following CLI command:

```
config bootconfig choice primary license-file <file>
```

OR

enter the following ACLI command:

```
(config)# boot config choice primary license-file <file>
```

 **Note:**

The variable '<file>' supports the following values for the source of a license file on an Ethernet Routing Switch 8800/8600:

- /flash/<file_name>
- /pcmcia/<file_name>
- <a.b.c.d>:<file_name>, where <a.b.c.d> is the IP address of an FTP or TFTP server

Considerations for upgrades from 5.0-based code releases

Users should read and reference the latest version of CSB 2008008618, Software Life-Cycle Management for the ERS 8800/8600 product, before deciding to move to any code release.

 **Important:**

For switch cluster systems running 5.0.0.x code (where x is less than 2), intermediate upgrades first to 5.0.0.2, then to one of 5.0.1.x, or 5.1.x are required, versus a direct upgrade to 7.1.0.0. 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. Refer to the 5.0.1.0 Release notes for details regarding the intermediate upgrade. Direct upgrades to release 7.1.0.0 are supported from 4.1.8.2, 4.1.8.3, 5.0.x (where x is 1 or higher), and 5.1.x.

Configuration file modifications for BGP upgrades from release 4.x code

 **Caution:**

Users using BGP with release 4.x code need to be aware of the following limitations regarding upgrading to 5.x or later code release. For any user using the add-as-path command in 4.x or earlier releases, a direct upgrade to 5.x or later code (including 5.0.0.x, 5.0.1.0, 5.1.0.0, 7.0.0.0, or 7.1.0.0 code) will create issues with your BGP operation, as the format for this command has changed in 5.x and all future code releases. The 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.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, if such config entries are in a 4.x config file, those entries must be manually converted to 5.x or later format before upgrading; the upgrade to 5.x or later 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 8800/8600 switch with its Unix VI editor), edited (for example with MS Word) and then copied back before upgrading.

For example, the above 4.x config example:

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

Must be changed to the following 5.x config format:

```
ip as-list 1 create 100 permit "64521"  
(Q01977204)
```

SMLT switch cluster upgrade considerations

With SMLT switch cluster upgrades, to maintain remote Telnet access to the switches, you must follow specific upgrade steps in some scenarios when upgrading to any higher release of code.

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, 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 5.1.0.0) to 7.1.0.0.




Figure 1: 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 7.1.0.0. Due to the SMLT behavior, the network diagram now looks like the following figure.




Figure 2: SMLT upgrade scenario after upgrading switch A to 7.1.0.0

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 7.1.0.0, 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 7.1.0.0 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 that you cannot use SSH in this upgrade scenario, as you cannot open SSH connections from one Ethernet Routing Switch 8800/8600 to another. You must use Telnet.

 **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 5.0.1.0 (or 5.1.0.0), and then to 7.1.0.0.

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 7.1, HA-CPU supports the following in Hot Standby mode:

- Shortest Path Bridging MAC (SPBM)
- platform configuration
- Layer 2 protocols: IGMP, STP, MLT, SMLT, ARP, LACP, VLACP
- Layer 3 protocols: RIP, OSPF, VRRP, RSMLT, VRF Lite

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

In Warm Standby mode, configuration synchronization is supported, but protocol state synchronization is not. Therefore, after failover, the protocols are restarted. These protocol restarts can result in small expected network down time.

HA-CPU supports the following in Warm Standby mode.

- DVMRP, PIM-SM, PIM-SSM
- BGP
- MPLS
- BFD
- IPv6, and all associated IPv6 protocols

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

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

PGM

Ongoing considerations

The following sections describe considerations that are not new for Release 7.1.0.0, but which still apply for 7.1.0.0.

Module and chassis compatibility and performance considerations

Release 7.1 does not support classic modules. Only R, RS, and 8800 series line card modules are supported with release 7.1. Also, the 8003 chassis is not supported with release 7.1. The 8003-R chassis replaces the 8003 chassis.

For switch fabric modules, only the 8692 with SuperMezz and 8895 CP/SF are supported with release 7.1.

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 6: 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 |
|----------|--|--|
| 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 |
| 8848GB | 16 Gbps | 60 Gbps |
| 8848GT | 16 Gbps | 60 Gbps |
| 8834XG | 16 Gbps | 60 Gbps |

If you place an R, RS, or 8800 module into a Standard slot of a non-high performance chassis, you receive the following message:

```
For maximum performance, Avaya 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 R, and RS modules.

Identify the High Performance Backplane by using the CLI or ACLI. Use the CLI command `show sys info` or the ACLI 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 8800/8600 Chassis, TT-0507501A* on the Avaya support Web site.

Table 7: Chassis hardware revision

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

| Chassis model | Hardware Revision | H/W Config |
|---------------|--|---------------|
| 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 8800/8600 chassis can have their existing chassis exchanged and reworked. Order service part number N0060024. An advanced replacement unit is provided.

Switch clustering topologies and interoperability with other products

When the Ethernet Routing Switch 8800/8600 is used with other Ethernet Routing Switch products, the switch clustering bridging, unicast routing, and multicast routing configurations vary with switch type. Avaya recommends that you use the supported topologies and features when you perform inter-product switch clustering. For more information, see *Switch Clustering Design Best Practices, NN48500-584* and *Large Campus Technical Solutions Guide, NN48500-575*, available on the Avaya support Web site.

SF/CPU protection and loop prevention compatibility

Avaya recommends several best-practice methods for loop prevention, especially in any Ethernet Routing Switch 8800/8600 Switch cluster environment. For more information about loop detection and compatibility for each software release, see *Large Campus Technical Solutions Guide, NN48500-575* and *Switch Clustering Design Best Practices, NN48500-584*.

Switch behavior during boot cycle and redundant configuration files

Avaya recommends that you take special care when providing the boot option for your production systems. The Ethernet Routing Switch 8800/8600 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 8800/8600 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 *Avaya Ethernet Routing Switch 8800/8600 Administration, NN46205-605*.

In normal operation, Avaya 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, Avaya 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.



Important:

If you want to store only one simple backup configuration file, Avaya 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.

This configuration works especially well with SMLT because of the other redundant switch in the SMLT cluster.

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 ACLI modes, unless otherwise specified.

Table 8: Switch behavior during boot cycle

| Parameters | Switch action |
|--|--|
| A configuration file is not specified. The config.cfg file is present on the flash drive. | The switch boots config.cfg |
| The primary configuration file is specified. The configuration file is present on the flash drive. | The switch boots the specified configuration file. |

| Parameters | Switch action |
|---|---|
| The primary configuration file is specified. The configuration file is not present on the flash drive. | The switch boots with factory defaults (if <code>config boot flags verify-config</code> is true , and a backup configuration is not specified). |
| The primary configuration file is specified. The configuration file on the flash drive has a bad command. | The switch boots with factory defaults (if <code>config boot flags verify-config</code> is true , and a backup configuration is not specified). |
| The primary configuration file is specified. The configuration file on the flash drive has a bad command. 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. In ACLI 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 ACLI 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.

1. To change the runtime configuration file locations, use the following command:

```
config bootconfig choice <primary|secondary|tertiary>
[config-file <file>|backup-config-file <file>|image-file
<file>]
```

For example, to specify the configuration file in flash memory as the primary, use the following command:

```
ERS-8610:6# config bootconfig choice primary config-file /
flash/config.cfg
```

2. To set the location for the I/O module driver image for the BootStrap protocol:

```
config bootconfig bootp image-name <image-name> <slot-number>

config bootconfig bootp secondary-image-name <image-name>
<slot-number>
```

For example, to specify an R module driver for slot 2 in flash memory, use the following command:

```
ERS-8610:6# config bootconfig bootp /flash/p80j50xx.dld 2
```

 **Important:**

Avaya recommends that you store .dld files in flash memory, and that you always set the image-name to default.

3. To set the boot source location for the SuperMezz image:

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

For example:

```
ERS-8610:6# config bootconfig mezz-image image-name /flash/
p80m50xx.img
```

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

1. Configure the primary configuration file choices:

```
config bootconfig choice primary config-file /flash/
primaryconfig.cfg
```

```
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
```

OSPF warning message

When you enable OSPF on a VLAN or a brouter port, if no OSPF area is associated with the interface (that is, the OSPF area for the interface is 0.0.0.0), the following warning message is displayed:

```
When enabling OSPF for a VLAN, this automatically creates area 0.0.0.0 for the
switch, which once the VLAN is active (VLAN has active ports) will result in the
advertisement of area 0.0.0.0 by this switch. If this is not the users intent, care
must be taken to place the VLAN into some other properly configured area. Area
0.0.0.0 will always be present for the switch, BUT this area will only be advertised
if some active VLAN exists and is assigned to area 0.0.0.0, which is the default
assignment.
```

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 8800/8600, 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, Avaya recommends that you increase the Hold Timer to 120 or 180 seconds.

IPv6 considerations

The switch cannot learn a given IPv6 neighbor's address on more than one interface (including link-locals). If the same address is learned on more than one interface, this can cause the switch to generate errors, such as:

```
swF:5# CPU5 [01/19/09 03:27:21] RCIP6 ERROR rcip6RpcOutChangeResEntrySubCid: | |  
REPLACE neighbor to HW FAILED. nbr ip address:
```

In a triangle SMLT, if you delete VRRP peers on the SMLT aggregation switches, the VRRP addresses on the data closet switch are not immediately cleaned up in the IPv6 neighbor table (`show ipv6 neighbor info`). The table shows IPv6 neighbor states as `Incomplete`. The neighbor addresses are only aged out 30 minutes after the traffic is stopped from the neighbor, in accordance with the ND RFC. In addition, the switch does not immediately delete router neighbors. Instead, it places them in the `Incomplete` state when they no longer exist. In this case, the virtual addresses are removed by the neighbor 30 minutes after deleting the VRRP virtual routers on the two switches.

SNMP considerations

SNMP is configured differently in the ACLI than in the CLI. Auto-generation of several parameters and command structure changes means that several configuration procedures are no longer required in the ACLI. These considerations only apply to upgrades from Release 4.x

to 7.1 as release 5.x already implements these changes. For more information, see the following:

- For SNMP trap changes, see the ACLI SNMP trap configuration section in *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting, NN46205-703*.
- For SNMP community-based changes, see *Avaya Ethernet Routing Switch 8800/8600 Administration, NN46205-605*.

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 7.1 does not support PIM Multicast Border Router (MBR) functionality over SMLT.

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

To improve SMLT failover and recovery behavior for large-scale networks, Avaya has optimized the IST protocol and rearchitected the SMLT state machines. This functionality improvement is mainly targeted for large-scale SMLT networks.

For best network operation, Avaya recommends that you operate switch clusters using only the new SMLT architecture. Within an SMLT cluster, you must run the same software release on both peer IST switches (except during upgrades).

The SMLT re-architecture is supported in releases 4.1.8.2, 4.1.8.3, 5.0.x (where x is 1 or higher), 5.1.x, 7.0.0.0, and 7.1.0.0.

In a scaled SMLT SPBM network environment, Avaya recommends increasing the aging timer from the default to 1 hour or more for VLANs.

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. To resolve this situation, you must bring the peer back online and save the configuration on both switches.

IST considerations

In EDM (or any SNMP based tool), whenever you change the MltType of an MLT to istMLT, configure the IST PeerIp and VlanId (1..4094) before you save the configuration. If you save the configuration without configuring the PeerIp and VlanId, you create an invalid configuration that cannot load during the booting process, which results in all the cards on the switch being taken off-line. (Q02132456)

60 day trial license

You are provided a 60 day trial period for the Ethernet Routing Switch 8800/8600, during which you have access to all features. In the trial period you can configure all features without restriction. The switch logs trial period expiration 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, the trial period expiration messages cease. At the end of the trial period, a message appears notifying the user that the trial period has expired.

After the license expires, configured licensed features are no longer functional after the switch is restarted or rebooted. If you want these configured features to continue to function properly, you must install a valid license.

For additional information about trial licenses, see *Avaya Ethernet Routing Switch 8800/8600 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 500ACEs 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 Avaya 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:  
ercdFilterRdxResultUpdate: ercdGetCollapseNode() Failed !!
```

MTBF for 1 Gig SFPs

The mean time between failure (MTBF) for all 1 Gig SFPs is 807 000 hours.

Supported standards, RFCs, and MIBs

For information about supported standards, RFCs, and MIBs, see the Appendices in *Avaya Ethernet Routing Switch 8800/8600 Planning and Engineering — Network Design, NN46205-200*.

Supported traps and notifications

For a complete list of log messages generated by Ethernet Routing Switch 8800/8600 Software Release 7.1, see *Avaya Ethernet Routing Switch 8800/8600 Logs Reference, NN46205-701*.

For a complete list of SNMP traps generated by Ethernet Routing Switch 8800/8600 Software Release 7.1, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting, NN46205-703*.

Chapter 4: Resolved issues

This section details all issues resolved for Release 7.1.

Platform resolved issues

Table 9: Platform resolved issues

| CR references | Description |
|---------------|---|
| Q02025261 | With multiple HA failovers, intermittent connectivity issues may occur. |
| Q02100062 | After a reboot of an HA switch, the following error may appear: CPU6 [12/10/09 02:32:53] COP-SW-IP ERROR Slot 4: ercdProcIpRecMsg: Failed to Add ECMP IP Record. IpAddr:xxx.xxx.xxx.xxx IpMask: xxx.xxx.xxx.xxx NumEcpRouteRecs: 1 retCode: 18 No traffic issues are seen with this message. |
| Q02126115 | If 8800/8600 switches running pre-7.0 code are connected to rebranded 8800 7.0 switches, the pre-7.0 switches cannot identify the chassis type and remote port from Topology Discovery Packets from the rebranded 8800 switches. As a result, in the pre-7.0 switches, the command show sys topology displays <code>unknown error: 192</code> in the ChassisType and Rem Port fields for the 8800 switches. |
| Q02132373 | With the 8895 SF/CPU, to copy files from either the master or secondary SF/CPU to an external device, do not use FTP or SCP, but rather use TFTP. If you use FTP or SCP to copy files from the SF/CPU, this action can lead to switch abnormalities. To copy files from an external device onto the 8895 SF/CPU, you can use TFTP, FTP, or SCP. |
| Q02135428 | There is a potential that the ERS 8800/8600 (8692 SF/CPU with Mezz only) can hang on boot when there is a version mismatch between the B and M images. If a switch reaches this state, a power cycle or reboot to fix the hang can lead to the <code>/flash</code> partition being reformatted. It is best practice to ensure that your image versions always match and that necessary files in the <code>/flash</code> partition are always backed up. |

Switch management resolved issues

Table 10: Switch management resolved issues

| CR references | Description |
|---------------|--|
| Q02091999 | The CLI can support a DNS host name of up to 256 characters; however, EDM can only support up to 64 characters. Therefore, do not configure a DNS host name greater than 64 characters. |
| Q02133713 | With the 8895 SF/CPU, the out-of-band management port now only operates with autonegotiation enabled. Autonegotiation cannot be disabled on the out-of-band management port. Further, for proper operation of the 8800 device, the 8895 management port must only be connected to a device that supports and is enabled for Autonegotiation and must also run in full duplex mode. Device connections that do not support autonegotiation and full duplex are not supported. |

KHI resolved issues

Table 11: KHI resolved issues

| CR | Description |
|-----------|---|
| Q02094865 | With the new KHI enhancements, the switch produces more messages at boot up. These messages only indicate issues if they appear concurrently with switch operational issues. |
| Q02102285 | KHI may report a false positive of bad packets on a 10-Gig port, even when nothing is plugged into the port. In addition, the error message refers to the lane rather than the port. For information about which ports are associated with which lanes, refer to Avaya Ethernet Routing Switch 8800/8600 Planning and Engineering — Network Design Guide (NN46205-200). |

Layer 2 resolved issues

Table 12: Layer 2 resolved issues

| CR references | Description |
|---------------|---|
| Q02053232-01 | RSTP/MSTP log messages introduced in 4.1.3.0 code and missing in 5.0.x and 5.1.x (but added in 5.1.2.0), are also missing in 7.0.0.0. |

MLT/SMLT resolved issues

Table 13: MLT/SMLT resolved issues

| CR references | Description |
|---------------|---|
| Q02099875 | In full mesh SMLT, if a VLAN IP is changed, the IST peer is not displaying the new IP address in the topology table. Workaround: save the config file where the IP address was changed and reboot the switch with that config file. After the switch comes back up, the IST peer will learn the new IP address with topology. |
| Q02102162 | The 8800 switch allows you to configure more than the supported maximum of 128 MLTs. Do not configure more than 128 MLTs on the switch as this is not a supported configuration. |
| Q02119996 | When you enter the <code>show ip rsmult info</code> command, the same SMLT ID can display twice for some VLANs. This issue arises only for VRF-enabled VLANs running RSMLT. |

Unicast routing resolved issues

Table 14: Unicast routing resolved issues

| CR references | Description |
|---------------|---|
| 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 |
|---------------|--|
| Q02089739 | <p>In a triangular SMLT setup with VRRP, if you delete the VRRP instance on the master router, the following error may appear:</p> <pre>*Dist-1-187:3# CPU3 [11/06/09 02:41:41] RCIP6 ERROR rcip6RpcOutChangeResEntryState: ify_arte lookup failed fe80:0:0:0:212:83ff:fe7c:2204 cid 16779277</pre> <p>There is no traffic impact from this issue.</p> |

Multicast routing resolved issues

Table 15: Multicast routing resolved issues

| CR references | Description |
|---------------|--|
| Q02076924 | In a PIM-SM network, if a single-attached multicast source is removed from the network, its entries are never removed from the mroute tables and the entries continue to be displayed under show ip mroute info . |
| Q02111397 | For the IGMP <code>stream-limit-max-streams</code> parameter, if the default value is changed, the new stored value appears incorrectly in the show commands and in the config. |

IPv6 resolved issues

Table 16: IPv6 resolved issues

| CR | Description |
|-----------|---|
| Q02122414 | In the ACLI, the ping datasize command supports the datasize range for IPv4 only: 16-4076. It should also support the expanded datasize range for IPv6 ping of 16-65487, as in the CLI. |
| Q02122417 | With IPv6, any ping executed with a data size above 1864 is dropped. Do not set the IPv6 ping data size above 1864. |
| Q02122887 | If you configure IPv6 VLANs, save the configuration, and then reboot, the offset used to create the VLAN MAC addresses can change, changing the VLAN MAC addresses, and in turn, the link-local IPv6 addresses. The link-local addresses can move from one VLAN to another. This can cause errors to appear in other network nodes such as the following: CPU5 [03/02/10 01:44:52] HW INFO replaceIpv6NbrRecordToBinTable: Unable to update |

| CR | Description |
|-----------|--|
| | neighbor record, a record with the same link-local address exists on a different interface This error can be accompanied by RCIP6 errors such as the following: CPU5 [03/02/10 01:44:52] RCIP6 ERROR rcip6RpcOutChangeResEntrySubCid: REPLACE neighbor to HW FAILED. nbr ip address: fe80:0:0:0:214:dff:fe52:265:, HAL error code = -1 Workaround: Administratively disabling and reenabling IPv6 on all of the VLANs on the node reporting errors clears the condition. Alternately, disabling and reenabling all the VLAN ports in IPv6 VLANs accomplishes the same. If it is known on which VLAN the link-local address was previously present and which VLAN it moved to, it is sufficient to disable and reenable IPv6 on these VLANs only. |
| Q02125173 | In ACLI mode, if you configure an IPv6 interface on a VLAN and then add the VLAN member ports to an MLT, the IPv6 interface goes down. This issue occurs only if the box is booted in MSTP mode RSTP. The issue does not happen if the box is booted in MSTP mode Default. As a workaround, after you add ports to the MLT, disable and reenable IPv6 on the VLAN. |

CLI and ACLI resolved issues

Table 17: CLI and NNCLI resolved issues

| CR references | Description |
|---------------|--|
| Q02087492 | With protocol-ID based VLANs, the DSAP/SSAP entry must be a four-digit hexadecimal number in the range 0x0 to 0xffff. If the first letter in the four-digit hexadecimal format entered is an invalid digit, an 'out of range' error is displayed. However, the CLI can accept four-digit numbers that include nonhexadecimal letters (for example: bbhh, 3dhj, abzz, and so on). If the trailing digits are invalid digits, the value is accepted and the valid part of the given number is extracted (for example: if bhhh is entered, the switch extracts a value of 0x000b). |
| Q02098992 | The <code>config ethernet <slot/port> fc-pause-0 <enable disable></code> command for configuring Ethernet ports does not apply to R/RS modules. This command will be removed in a future release. |
| Q02100377 | In the CLI, you can configure a DNS host name of up to 256 characters. The CLI should limit the DNS host name to 64 characters. |
| Q02100686 | In a certain case, SMLT entries can occur twice in the ACLI config file. First, the switch must be in STP mode. Then, you must save the config to ACLI |

| CR references | Description |
|---------------|--|
| | and then boot that config. Finally, if you configure an SLT port and add the SMLT ID to that port, the port will have two SMLT entries. The additional lines in the config file have no effect on normal switch operations. |
| Q02101603 | When running in HA mode, the login prompt may scroll on the console screen. This issue does not appear if the CLI timeout is set to the default value of 900. As a workaround, restore the CLI timeout to the default value of 900 (using CLI or EDM [Security > control path > general > CLI]). |
| Q02117793 | Sourcing of large ACLI config files can take several minutes, while at booting takes only 10 seconds. The issue is seen while sourcing at runtime only. Booting is normal. |
| Q02121585 | In the ACLI, with IP-subnet based VLANs in MSTP mode, do not configure a name for the VLAN otherwise the saved configuration can cause issues. |
| Q02124930 | In ACLI mode, the show fulltech command displays the chassis type as 8810co rather than 8010co. |

Enterprise Device Manager resolved issues

Table 18: Enterprise Device Manager known issues

| CR | Description |
|-----------|--|
| Q02076555 | In EDM, under IP > BGP > Aggregates , if you modify the parameters of an existing Aggregate entry (for example AsSetGenerate, SummaryOnly, SuppressPolicy, AdvertisePolicy, and AttributePolicy) and click Apply , the change is not displayed accordingly in some cases. To display the correct values, refresh the EDM screen. This issue can also occur when you delete an Aggregate entry. |
| Q02077395 | With EDM, if you attempt to delete forwarding database (FDB) entries using the Forwarding tab under Configuration > VLAN > VLANs , an error is displayed and the FDB entries are not deleted. To delete FDB entries in EDM, use the Configuration > VLAN > VLANs > Advanced tab. For the desired VLAN, double-click the VlanOperationAction table cell, select flashMacFdb from the drop-down menu and click Apply . You can also use the config vlan <vid> fdb-entry flush command (in the CLI) or the vlan mac-address-entry <vid> flush command (in the VLAN Interface Configuration mode of the ACLI). |
| Q02088297 | In EDM, after a port is assigned to a VRF, the user can manage this port from the assigned VRF including creating an IP Router port, OSPF |

| CR | Description |
|-----------|---|
| | interface, and a RIP interface. From the GRT, the user can manage the basic functionality for the port, for example disabling and enabling the port, but cannot manage the IP functionality for the port. If the user configures an IP address on a port from the VRF, the GRT cannot display this data, and no IP functionality on this port can be managed from the GRT. Due to this problem, EDM shows no data or wrong values in the GRT from the Edit > Port > IP path. |
| Q02089610 | EDM allows you to configure the IPv6 OSPF stubmetric parameter (under IPv6 > OSPF > Areas) beyond the valid range of 0-65535 without producing an error. |
| Q02091957 | Under Security > Datapath > ACL filters > ACL > ACE , if you select an existing ACE and click Action/Debug , and then click the ellipsis (...) to select a DstMltId from the pop-up window, to remove the selected value, you must deselect the DstMlt value and also deselect either the associated DstPortList or the DstVlanId to remove it as well. |
| Q02097130 | Under Edit > Diagnostics > PCAP > PcapGlobal , if you modify the BufferSize field to a value that consumes too much memory, an <code>UndoFailed</code> error is displayed. EDM should display an error similar to the following: <code>Possible Memory allocation failure - please refer to logs in PCAP engine</code> |
| Q02099531 | In EDM, under the IP > Policy > Route Policy tab, if you double-click the MatchAsPath or MatchCommunity fields, values are duplicated in the pop-up window. If you assign one of the duplicate values to the MatchAsPath or MatchCommunity field, it gets applied. Once applied, do not attempt to assign the other duplicate value to the MatchAsPath or MatchCommunity field; otherwise an error is displayed. |
| Q02109487 | In this release, EDM help is unavailable for the EDM Quick Start pages (Configuration > Quick Start > Quick Start). Therefore, in these pages, the Help button is disabled. |
| Q02122686 | In Release 5.1 Java Device Manager, the UpdateSourceInterface field was a configurable option under IP > BGP > Peers . In EDM, under the IP > BGP > Peers tab, the same field is unavailable. As a workaround, you can use the <code>config ip bgp neighbor update-source-interface</code> (CLI) or <code>ip bgp neighbor update-source</code> (ACLI) commands to configure the field. |

Off-box EDM plug-in resolved issues

Table 19: Off-box EDM plug-in resolved issues

| CR | Description |
|-----------|--|
| Q02127403 | In the off-box EDM plug-in, the following issue can occur with multiple port configuration of NSNA. First, select multiple ports, then right click and select Edit General , and then click the NSNA tab. From the NSNA tab, if you set the mode to dynamic or uplink, and then attempt to select UplinkVlans or VoipVlans by clicking the associated ellipsis (...), an empty box is displayed. To work around this issue, you can configure NSNA on each port individually. Related to the above, if you set the mode to disabled, the UplinkVlans or VoipVlans fields should be greyed out. |
| Q02127410 | In the off-box EDM plug-in, the following issue can occur with multiple port configuration of FDB protect. If you select multiple ports, then right click and select Edit General , then click the Fdb Protect tab, none of the data is displayed for this tab. To work around this issue, you can configure FDB Protect on each port individually. |

Chapter 5: 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 7.1 known issues](#) on page 67
- [Previously reported known issues](#) on page 69

Release 7.1 known issues

Table 20: Release 7.1 known issues

| WI references | Description |
|---------------|---|
| wi00564305 | Enabling MSTP on IST ports enables the forceportstate of the IST ports internally. If you toggle the forceportstate with the forceportstate <enable disable> command, you will get a consistency error message. Do not disable forceportstate of IST ports once Spanning Tree is enabled on the IST. |
| wi00564333 | When the CP Limit is set to 100000 packets per second and the broadcast or multicast traffic across the port exceeds that limit, the port does not shut down as expected. |
| wi00700896 | KHI error messages are sometimes displayed after rebooting a switch in ACLI mode. These messages appear intermittently and do not cause any traffic loss. |
| wi00824067 | On reboot of the 8895 SF/CPU, the following message appears: <pre>SWA_7000-slot6:0x51aa300 (ttNetTask): mBlkClFree -- Invalid mBlk</pre> This is an intermittent error message that can be safely ignored. |
| wi00824070 | After rebooting a switch multiple times, there is a possibility that the switch can hang. |

| WI references | Description |
|---------------|---|
| wi00824072 | After a switch boot, if you attempt to execute the format /pcmcia command, console/telnet hangs. |
| wi00825446 | After upgrading a system with 8692 SF/CPU with Supermezz and HA to release 7.1, multiple reboots occur on the SF/CPU during the boot process before the SF/CPU boot successfully. This problem is intermittent and only occurs when booting the nodes from images residing on the /pcmcia file system. |
| wi00833969 | In a ring topology with OSPF and IS-IS configured in the core, a core link break causes slow convergence that may lead to SPBM L2 traffic loss. If the last member link of an OSPF VLAN fails, it takes down the IP interface and OSPF has to reconverge. While OSPF is reconverging, SPBM cannot get any CPU time so there is some traffic loss. |
| wi00841377 | When you click on the Help button for the LSP Summary tab, an incorrect Help page displays. The correct page should display the Displaying LSP summary information procedure, which you can see in <i>Avaya Ethernet Routing Switch 8800/8600 Configuration — Service Provider Bridging MAC (SPBM)</i> (NN46205–525). |
| wi00850194 | <p>When using VPFM discovery or doing a MIB walk on a scaled device, the query of MIB dot1dTpFdbTable could take a very long time to find the next entry. The CPU usage could spike to 100% during this time. Workaround: Disable the MIB dot1d-tp-fdb-query that takes a long time to get the next entry.</p> <p> Note: Keep in mind that you will not have access to this MIB after it is disabled. CLI command for the workaround: config sys set dot1d-tp-fdb-query disable ACL command for the workaround: no sys dot1d-tp-fdb-query</p> |
| wi00853802 | IS-IS uses TLV 135 to propagate routing information. This TLV is not user configurable so you cannot modify the metric. |
| wi00853908 | VLACP goes down if you delete the last C-VLAN on an NNI link and a default-vlan-id is set on the port. Workaround: Set the default-vlan-id to 0 before deleting the last C-VLAN or delete the default VLAN ID last. |
| wi00854252 | On the 8692 SF/CPU module with SuperMezz, the reset command boots the switch with a warm boot instead of a cold boot. Workaround: Use the following steps to reset with a cold boot: <ol style="list-style-type: none"> 1. Enter the reset command to reset the switch. 2. Stop the switch in boot monitor mode before loading. 3. Enter the reset command in boot monitor mode. |

| WI references | Description |
|---------------|--|
| wi00855097 | Do not configure VLAN 4093. This VLAN was for a legacy module that is no longer supported. If you configure this VLAN, it will not work as expected. |
| wi00855106 | FDB filters for PCAP were for legacy modules <i>only</i> . This feature is not supported on Release 7.1 R and RS modules. |
| wi00857202 | There are two show port roles commands: one displays roles per <i>vlan</i> and one per <i>portlist</i> . The show ports info rstp role vlan <vid> command does not work. However, the show ports info rstp role port <portlist> command works as expected. |
| wi00857629 | <p>There are discrepancies between the CLI and EDM in how you configure BGP 4-byte AS when the as-dot flag is enabled.</p> <ul style="list-style-type: none"> • The CLI allows you to enter the AS number in plain format or in as-dot format. EDM allows you to enter the AS number in AS-dot format only. • The CLI accepts a blank value for an AS number and assigns the default value of zero. For example, if instead of entering "x.y" you only enter "x. ", the CLI assigns 0 for the "y" value (x.0). EDM requires both an "x" and "y" value. |

Previously reported known issues

The following sections list known issues in Ethernet Routing Switch 8800/8600 reported in software releases prior to Release 7.1. These are to be resolved in a future release.

Platform known issues

Table 21: Platform known issues

| CR references | Description |
|---------------|---|
| wi00506367 | Line RDI is not generated properly as a result of LOS on the 8683XZR module in WAN mode. |
| wi00506474 | 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. |
| wi00506722 | 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. |

| CR references | Description |
|---------------|---|
| wi00507101 | On 8612XLRS modules with DDM enabled, wait 3 minutes after module initialization before you enter show sys pluggable-optical-modules commands to avoid errors during the initialization. |
| wi00507117 | 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. |
| wi00507119 | After a reboot, a COP software error message similar to the following may be displayed on the switch: CPU5 [10/30/09 11:23:06] COP-SW ERROR 27806496: LtrId = 152,LtrPrio=0,ltrStatus=15(LTR_SYNC_MSG_SLOT_INUSE),msgId=152,msgState = 1,Slot=4 You can ignore this message as it does not cause any functional issues. |
| wi00517103 | Configuring Distributed MLT on ERS 8800/8600 using 8608GBE cards could lead to a ports being put in different STP states in the MLT. |
| wi00517507 | Ping does not work when the source IP option is set to a circuitless IP interface. |
| wi00517523 | 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. |
| wi00517565 | With DWDM XFPs, the show system pluggable-optical-modules threshold status is incorrect during transition from a "High Alarm" to a "Low Alarm". 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. |
| wi00517636 | 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. |
| wi00517817 | 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. |
| wi00518502 | On reboot of the 8895 SF/CPU, the following message appears: <pre>SWA_7000-slot6:0x51aa300 (ttNetTask): mBlkClFree -- Invalid mBlk</pre> This is an intermittent error message that can be safely ignored. |

| CR references | Description |
|---------------|--|
| wi00518565 | When upgrading FPGA firmware on R or RS modules, the following message can appear: <code>Router-C:5#/CPU5 [03/08/10 15:04:15] COP-SW ERROR 27894800: LtrId = 152,LtrPrio=0,ltrStatus = 15 (LTR_SYNC_MSG_SLOT_INUSE),msgId=53,msgState =1,Slot=3</code> This message has no negative effect on the FPGA upgrade. There are no specific FPGA upgrades required with release 7.0. |
| wi00518632 | Some users prefer to copy files to and from the flash using Windows instead of using TFTP or FTP. However, release 7.0 does not recognize flash files formatted with either FAT16 or FAT32. As a workaround, you can TFTP or FTP the files to the flash. |
| wi00518661 | If you enable DDM monitoring on a switch with non-DDM GBICs installed, the switch generates a message (<code>HAL INFO GBIC</code>) every 5 seconds to the console and to the log file for each non-DDM GBIC installed. |
| wi00518690 | When rebooting the master CPU, the following warning messages may appear on the 8895 SF/CPU: <code>nyhq-csbu-udb:6# 0x51bb4e0 (tNetTask): duplicate IP address 2f50ef10 sent from ethernet address f4:e3:b1:0d:14:00</code> In addition, the following messages may appear on either the 8692 or 8895 SF/CPU: <code>CPU6 [11/24/09 12:15:49] MLT WARNING smltTick: pollCount = 51 > 50. But IST Channel active and resetCount = 0 < 3. Resetting pollCount and staying active!. CPU5 [11/24/09 12:17:44] IP INFO the Rsmlt circuit of vlan 18 is existed already in slave CPU No traffic issues are seen with these messages.</code> |
| wi00518696 | For the system power supply calculation, a low inaccurate value (410 W) is associated with any power supply that displays as <code>unrecognized</code> . This can lead to a system power calculation stating the system does not have enough power, when in fact it does. Properly installed Avaya-manufactured power supplies do not display as <code>unrecognized</code> . |

Switch management known issues

Table 22: Switch management known issues

| CR references | Description |
|---------------|--|
| wi00517339 | When configuring SSH on the switch, <code>-C</code> and <code>-C2</code> compression options are accepted, but should be rejected. Subsequent SSH connection are |

| CR references | Description |
|---------------|---|
| | also accepted with no message to the user. The switch should prompt the user with a message stating compression is not supported. |

KHI known issues

Table 23: KHI known issues

| CR | Description |
|------------|--|
| wi00508040 | <p>If you reset a slot that is passing traffic, the following KHI errors can result:</p> <pre> :5# CPU5 [11/04/09 06:52:01] KHI WARNING Port 4/2 is experiencing Packet Errors :5# CPU5 [11/04/09 06:52:01] KHI WARNING Port 4/4 is experiencing Packet Errors, Frames Long Errors :5# CPU5 [11/04/09 06:52:13] KHI WARNING Port 4/6 is experiencing Packet Errors, FCS Errors :5# CPU5 [11/04/09 06:52:24] KHI WARNING Slot 4 Middle Lane is experiencing Ingress RSP Errors - PM EME1 Parity Error :5# show bootconfig CPU5 [11/04/09 06:53:48] KHI WARNING Slot 4 Middle Lane is experiencing Ingress RSP AM Short Packets :5# CPU5 [11/04/09 06:53:48] KHI WARNING Slot 4 Middle Lane is experiencing F2X Errors - F2I Ingress SPI-4.2 Abort Received </pre> |

Layer 2 known issues

Table 24: Layer 2 known issues

| CR | Description |
|------------|---|
| wi00506797 | <p>If you disable a member port of an MLT that is running RSTP and then display statistics for the disabled port (for example, using the show spanning-tree rstp port statistics <slot/port> CLI command), the command output indicates that the port is still sending and receiving BPDUs. This is the normal display behavior for MLT ports. When the system displays the RSTP statistics for MLT ports, the statistics are taken from the designated port and displayed for all member ports. Even if a port is disabled, it is still a member of the MLT and hence the designated port's statistics are displayed for the disabled port. However, there are actually no packets going out the disabled port.</p> |
| wi00508462 | <p>In some cases, the output of the show slpp interface gig command does not show anything on either IST peer even when SLPP</p> |

| CR | Description |
|----|--|
| | has brought the ports down. The command should normally display some information on either one or both of the IST peers. |

MLT/SMLT known issues

Table 25: MLT/SMLT known issues

| CR references | Description |
|---------------|--|
| wi00517880 | For an IP VPN-lite configuration, where an edge 8800/8600 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. |
| wi00518481 | After rebooting an IST switch, the following error message may appear on the IST peer: <code>COP-SW-IPV6 ERROR Slot 7: ercdDeleteIPv6Record: Failed to lookup entry in gIPv6RadixTbl. Status: 18</code> There are no functional impacts from this issue. |
| wi00523290 | Consider a triangle SMLT network where the edge switch is connected to each IST switch using SMLT over MLT. On either of the IST switches, if you delete and re-add the SMLT interface to the edge switch, duplicate traffic to the edge switch can result. Workaround: 1. Before re-creating the SMLT on the MLT interface, shut down the ports of the MLT, and reenable them after assigning the SMLT ID. 2. After deleting the SMLT-ID, delete and re-create the MLT. |

Unicast routing known issues

Table 26: Unicast routing known issues

| CR references | Description |
|---------------|---|
| wi00505890 | 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. |
| wi00517472 | In OSPF Router LSA updates, the V-bit is not set, and is always 0. |
| wi00517787 | 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 |

| CR references | Description |
|---------------|---|
| | may show ExStart state for 5 to 8 minutes. The condition does clear itself in that time frame, and will go to full adjacency. |

Multicast routing known issues

Table 27: Multicast routing known issues

| CR references | Description |
|---------------|---|
| wi00506569 | SSM channel set to false on receiving traffic |
| wi00507488 | If MVR is enabled on the global level for a particular VRF and a configuration save is performed, the "MVR ENABLE" command is repeated two times in the configuration file. If the MVR is disabled after being enabled for the particular VRF, the configuration file shows the "MVR DISABLE" command followed by the "MVR ENABLE" command. This is done intentionally and is required for proper functioning of MVR on the Ethernet Routing Switch. The first "MVR ENABLE" command for a particular VRF does the job of allocating memory for all the structures required by MVR to run on the concerned VRF. The subsequent "MVR DISABLE" or "MVR ENABLE" command does the job of disabling or enabling the MVR feature on the VRF. The memory for MVR structures is never de-allocated unless the VRF is deleted or the switch is rebooted. Please do not edit the configuration file and delete either the "MVR DISABLE" or "MVR ENABLE" command considering them duplicate or redundant. |
| wi00507736 | The 8800 switch drops multicast traffic with source IP address of 0.0.0.0. |
| wi00518161 | RPF checks fail with MSDP peer configured in iBGP configuration. Workaround: Use default Peer or do not use iBGP configuration. |

IPv6 known issues

Table 28: IPv6 known issues

| CR | Description |
|------------|--|
| wi00507966 | In ACLI mode, if you configure an IPv6 interface on a VLAN and then add the VLAN member ports to an MLT, the IPv6 interface goes down. This issue occurs only if the box is booted in MSTP mode RSTP. The issue does not happen if the box is booted in MSTP mode Default. As a workaround, after you add ports to the MLT, disable and reenable IPv6 on the VLAN. |

CLI and ACLI known issues

Table 29: CLI and ACLI known issues

| CR references | Description |
|---------------|---|
| wi00506209 | The <code>copy</code> command does not work properly with FTP debug turned on. |
| wi00506337 | When sending traps to an Element Manager, the switch only uses the IP address specified by the first entered <code>sender-ip <dest-ip> <source-ip></code> command. It is possible to specify multiple sender IPs and each should use a different IP as specified in this command. The switch uses the IP address of the physical VLAN of the first entry in the target-address table. |
| wi00517661 | 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: <code>Another show or save in progress. Please try the command later.</code> |

Enterprise Device Manager known issues

Table 30: Enterprise Device Manager known issues

| CR | Description |
|---------------------------------|--|
| wi00507915 and wi00853479 | <p>IGMP SNOOP warning messages are not displayed in EDM. In EDM, go to VLAN > VLANs > IP > IGMP, check the SsmSnoopEnable checkbox, and click Apply. Because Snoop is not yet enabled, you should see the following message:</p> <pre>WARNING: IGMP SNOOP should also be enabled with IGMP SSM-SNOOP.</pre> <p>No such warning message is displayed. Using the CLI, make sure that both Snoop and SsmSnoop are enabled. Then in EDM go to VLAN > VLANs > IP > IGMP. Now try to uncheck SnoopEnable and click Apply. Because SsmSnoop is still enabled, you should see the following message:</p> <pre>WARNING: IGMP SSM-SNOOP should also be disabled with IGMP SNOOP.</pre> <p>No such warning message is displayed.</p> |
| wi00518024 | If you use EDM to export MLT configuration data (using the MultiLink/LACP Trunks tab, LACP tab, or IST/SMLT Stats tab under Configuration |

| CR | Description |
|------------|--|
| | <p>> VLANs > MLT/LACP), the display of the exported data is misaligned with the table header row. Although the data display is misaligned, the data values are correct.</p> |
| wi00518427 | <p>In the ACE Common EDM tab (under Configuration > Security > DataPath > ACL Filters > ACL > ACE), to configure the RedirectNextHopIpv6 parameter, you must first verify that the PktType field for the corresponding ACL (under Configuration > Security > DataPath > ACL Filters > ACL) shows IPv6. If the ACL is configured for IPv4, then the RedirectNextHopIpv6 configuration does not take effect. If you do configure the RedirectNextHopIPv6 field on an IPv4 ACL, while the IPv6 value is not saved, the RedirectNextHop field (for IPv4) can be populated with an erroneous IPv4 address. Be sure to delete the erroneous IPv4 address.</p> |
| wi00518439 | <p>In the EDM Physical Device view, EDM does not display the name of the 8692 SF/CPU cards. This issue does not affect 8895 SF/CPU cards.</p> |
| wi00518602 | <p>In EDM, if you set the VRRP FasterAdvInterval parameter (under Configuration > IP > VRRP > Interface) to a value that is not a multiple of 200 ms, no warning is displayed. A message similar to the following from the CLI should appear:</p> <pre data-bbox="493 911 1343 989">WARNING: Input value is not a multiple of 200ms, Fast Adv Interval adjusted to 200ms.</pre> <p>The warning is displayed if you modify the FasterAdvInterval under Configuration > VLAN > VLANs > IP > VRRP.</p> |
| wi00518694 | <p>In EDM, under Configuration > IP > DVMRP > Interface Advance, if you double-click the InPolicy or OutPolicy parameter and then click the Refresh button, the displayed pop-up window disappears. One workaround is to keep the lower scroll bar to the left-most position, in which case the refresh works and the popup window does not disappear.</p> |
| wi00518706 | <p>In EDM, if you create a BGP Peer (under Configuration > IP > BGP > Peers > Insert), the AdvertisementInterval value defaults to 30. This value should default to 5, which is the default route advertisement interval value for configuration using the CLI or ACLI.</p> |
| wi00518720 | <p>If you launch on-box EDM using Internet Explorer and then graph a port, you cannot change the default 5s polling interval from the drop down box. As a workaround, you can launch on-box EDM using Firefox, or use the off-box EDM plug-in.</p> |
| wi00523304 | <p>The work flow for creating an IP VPN route target in EDM differs from that for the CLI or ACLI. If you want to create a route target through EDM, you must perform the following steps:</p> <ol style="list-style-type: none"> 1. Select IP > IPVPN > Route Target > Insert. 2. Enter a valid index and IP address in the respective fields. 3. Click Insert. |

| CR | Description |
|----|--|
| | <ol style="list-style-type: none"><li data-bbox="508 247 1036 279">4. Select IP > VRF > Insert to create a VRF.<li data-bbox="508 296 1357 359">5. Select IP > IPVPN > VPN > Insert to create an IPVPN for the VRF just created.<li data-bbox="508 375 1357 470">6. For the IPVPN just created, change the importRTList or exportRTList to associate the route target (put the route target index for importRTList or exportRTList) with the IPVPN. |

Chapter 6: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com/support, or go to one of the pages listed in the following sections.

Navigation

- [Getting technical documentation](#) on page 79
- [Getting product training](#) on page 79
- [Getting help from a distributor or reseller](#) on page 79
- [Getting technical support from the Avaya Web site](#) on page 80
- [Additional information](#)

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

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

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.