



Release Notes — Software Release 7.2.10 Avaya Ethernet Routing Switch 8800/8600

Release 7.2.10
NN46205-402
Issue 11.08
March 2014

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 may generally make available to users of its products and Hosted Services. 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 websites 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 Avaya hardware and software. 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 website: <http://support.avaya.com> or such successor site as designated by Avaya. Please note that if you acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to you by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO](http://support.avaya.com/licenseinfo) OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER; 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").

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users.

Licence types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

"Heritage Nortel Software" means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at <http://support.avaya.com/LicenseInfo/> under the link "Heritage Nortel Products", or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants Customer a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or (in the event the applicable Documentation permits installation on non-Avaya equipment) for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, 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

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those Products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the Documentation or on Avaya's website at: <http://support.avaya.com/Copyright> or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components

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 website: <http://support.avaya.com> or such successor site as designated by Avaya. 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, Hosted Service(s), 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, Hosted Service(s) 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. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: <http://support.avaya.com>, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: <http://support.avaya.com> for Product or Hosted Service notices and articles, or to report a problem with your Avaya Product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: <http://support.avaya.com> (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Chapter 1: Introduction

Purpose of this document

This document describes important notices and fixed and known issues for Avaya Ethernet Routing Switch 8800/8600 Release 7.2.10. software.

Ethernet Routing Switch 8800/8600 Release 7.2.10 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 7.2.10 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 13) and unsupported ([Unsupported hardware for Release 7.2](#) on page 18) combinations of hardware and software, as well as new feature descriptions.

Related resources

Documentation

See the *Avaya Ethernet Routing Switch 8800/8600 Documentation Roadmap*, NN46205-103, for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at <http://avaya-learning.com/>.

Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

- To find videos on the Avaya Support website, go to <http://support.avaya.com>, select the product name, and check the *videos* checkbox to see a list of available videos.
- To find the Avaya Mentor videos on YouTube, go to <http://www.youtube.com/AvayaMentor> and perform one of the following actions:
 - Enter a key word or key words in the Search Channel to search for a specific product or topic.
 - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the site.

 **Note:**

Videos are not available for all products.

Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support web site.

About this task

You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), that apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Ethernet Routing Switch 8800.

Procedure


1. In an Internet browser, go to <https://support.avaya.com>
2. Type your username and password, and then click **LOG IN**.
3. Click **MY PROFILE**.




4. On the site toolbar, click your name, and then select **E Notifications**.



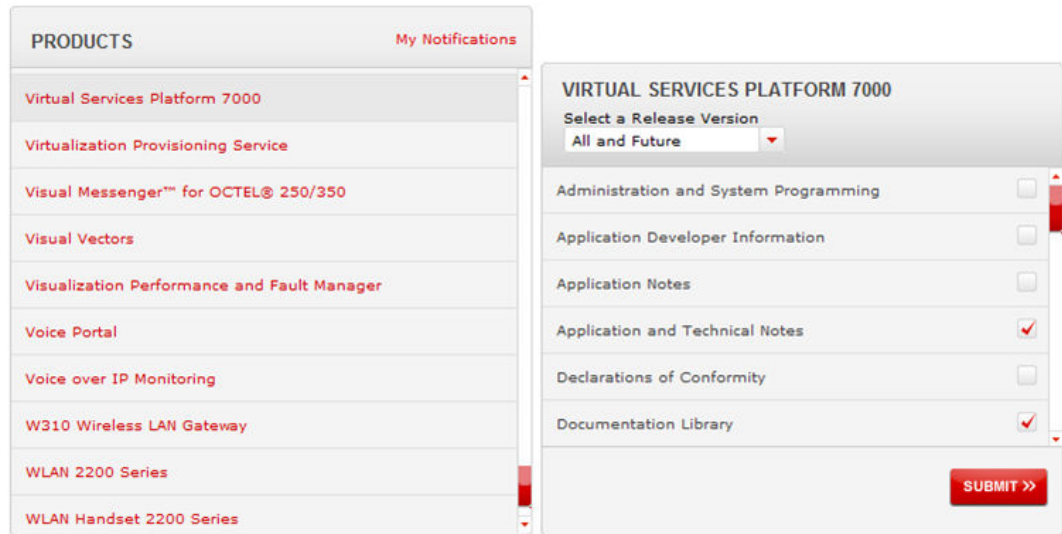
5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.



6. Click **OK**.
7. In the PRODUCT NOTIFICATIONS area, click **Add More Products**.



- 8. Scroll through the list, and then select the product name.
- 9. Select a release version.
- 10. Select the check box next to the required documentation types.



- 11. Click **Submit**.

Support

Visit the Avaya Support website at <http://support.avaya.com> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Chapter 2: New features in this release

This section describe the new features for the Avaya Ethernet Routing Switch 8800/8600 Release 7.2.10.

Bidirectional Forwarding Detection over IPv6

The Bidirectional Forwarding Detection (BFD) protocol detects forwarding path failures between two connected forwarding engines in sub-second time. The ERS 8600/8800 already supports BFD for IPv4 interfaces. In Release 7.2.10, Avaya extends BFD support to IPv6 interfaces. As part of this feature, OSPFv3 and Static Routing can register with BFD for quicker failure notifications.

For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IPv6 Routing* (NN46205–504).

IPv6 traceroute

In Release 7.2.10, the **traceroute** command supports IPv6 addresses. For more information, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting* (NN46205-703).

MLT hashing enhancement

Release 7.2.10 enhances the MLT hashing algorithm for IPv6 traffic by including more fields from the IPv6 header to achieve better hashing distribution. This includes both native IPv6 traffic and traffic migrating from IPv4 to IPv6 (sometimes called *6in4 tunneled traffic*). For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Link Aggregation, MLT, and SMLT* (NN46205-518).

LACP enhancement

In directly-connected physical links, LACP must be enabled globally and at the port level on both ends of the link.

Starting with Release 7.2.10, if LACP is not enabled on both ends, a Warning message informs the user that there is a port level or global level LACP link configured incorrectly. The message displays after a 180 second interval and not more than one LACP Rx past the 180 second interval.

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

Service Level Agreement (SLA) Monitor

SLA monitors help network administrators ensure that the network maintains a high level of data and voice communication quality by measuring network performance, monitoring IP services, and assisting with network troubleshooting.

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

Ping enhancement

Release 7.2.10 enhances the `ping` command by adding interface options for ports, management port, VLANs, and tunnels. For configuration information, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting* (NN46205-703).

Licensing enhancements

Release 7.2.10 includes the following licensing enhancements.

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

Warning messages

The ERS 8800/8600 switch generated a Warning message every 24 hours to notify customers that their feature license was not installed or the trial period expired.

This enhancement stops the display of erroneous messages by checking to see if the Advanced or Premier feature is actually configured on the switch before generating the Warning message.

VRF security risk

If a VRF Premier license was not installed or the trial period expired, there was a potential security risk because the VLANs of VRFs were added to the Global Router (Vrf 0).

This enhancement allows the VRF configurations to load properly, but you cannot enable or create any new Advanced or Premier features configurations until you install a valid license.

PCAP and port mirroring

If a PCAP Advanced license trial period expired, there was no way to remove the PCAP configuration. Because the PCAP filters could not be removed, the switch blocked any new port mirroring configurations.

This enhancement generates a log messages every hour indicating that the PCAP license is invalid. However, it also enables you to delete the PCAP configuration if you no longer want to use that feature. This enables you to continue to use port mirroring.

Network Time Protocol (NTP)

In previous releases of NTP, the source IP address was always the address of the outgoing interface. Starting with Release 7.2.10, you can specify a circuitless IP (CLIP) IP, a Management Virtual IP, or continue using the outgoing IP (default).

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

Link Layer Discovery Protocol (LLDP)

Link Layer Discovery Protocol (LLDP) enables stations connected to a LAN to advertise their capabilities to each other, enabling the discovery of physical topology information for network management. LLDP-compatible stations can consist of any interconnection device including PCs, switches, and routers. Each LLDP station stores LLDP information in a standard Management Information Base (MIB), making it possible for a network management system (NMS) or application to access the information.

Each LLDP station:

- advertises connectivity and management information about the local station to adjacent stations on the same 802 LAN (802.3 Ethernet with 8300 Series).
- receives network management information from adjacent stations on the same LAN.

LLDP also makes it possible to discover certain configuration inconsistencies or malfunctions that can result in impaired communications at higher layers. For example, it can be used to discover duplex mismatches between a router and the connected switch.

LLDP is compatible with IETF PROTO MIB (IETF RFC 2922).

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

New SFP+ devices for Release 7.2.10

Release 7.2.10 adds support for the following SFP+ devices:

- 10GBASE ZR/ZW SFP+ — This is a 1550 nm single-mode fiber SFP+ with a range of 70 km. The part number is AA1403016–E6.
- 10GBASE-ER CWDM SFP+ — These SFP+ devices have a reach up to 40 km. The following list shows the part numbers for devices with various wavelengths:
 - AA1403153-E6 — 1470 nm
 - AA1403154-E6 — 1490 nm
 - AA1403155-E6 — 1510 nm
 - AA1403156-E6 — 1530 nm
 - AA1403157-E6 — 1550 nm
 - AA1403158-E6 — 1570 nm
 - AA1403159-E6 — 1590 nm
 - AA1403160-E6 — 1610 nm

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 7.2.10, and provides important information for this release. Sections that specifically apply to 7.2.10 are labeled as such. Otherwise, the generic sections that apply to all 7.2.x releases are labeled as 7.2.

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		Introductory 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
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
8005AC	1462 W AC	4.0.0	DS1405012
8005DI AC	1462 W Dual input AC	5.0	DS1405018-E6

Item		Introductory version	Part number
Chassis			
8005DI DC	1462 W Dual input DC	5.1	DS1405017-E5
8005DC	1462 W DC	4.0.x	DS1405011

Compact flash support on 8895 SF/CPU

Avaya recommends using only the Compact Flash cards listed below with the 8895 SF/CPU since they have been validated for proper operation. Use of any other Compact Flash devices is not recommended as they have not been verified for compatibility on the 8895 SF/CPU.

- SSD-C02G-4000
- SSD-C02G-4007
- SSD-C02G-4300
- SSD-C02G-4500
- SSD-C02G-4600

Module and component compatibility

This section has three tables:

1. The first table lists all the Ethernet 8800, RS, and R modules supported in this release.
2. The second table lists the supported SFP/SFP+ transceivers.
3. The third table maps the modules to the transceivers that they support.

Table 2: Supported modules

Modules		Minimum software version	Part number
Ethernet 8800 modules			
8812XL	12 port 10 Gigabit Ethernet SFP+	7.1.3.0	DS1404121-E6
8834XG	24 100/1000 Mbps SFP ports 2 XFP ports 8 10/100/1000 Mbps copper ports	7.1.0.0	DS1404123-E6
8848GB	48 100/1000 Mbps SFP ports	7.1.0.0	DS1404122-E6

Modules		Minimum software version	Part number
8848GT	48-port 10/100/1000 Mbps copper ports	7.1.0.0	DS1404124-E6
Ethernet RS modules			
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
8648GTRS	48-port 10/100/1000 Mbps copper ports	5.0.0	DS1404110-E6
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

Table 3: Supported transceivers

Transceivers		Minimum software version	Part number
100BASE Small form factor pluggable transceivers			
100BASE-FX SFP	1310 nm LC connector	5.0.0	AA1419074-E6
1000BASE 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

Transceivers		Minimum software version	Part number
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
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 (XFP) 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
10 Gigabit Ethernet Small form factor pluggable (SFP+) transceivers			
10GBASE-LR	1310 nm SMF with a range up to 10 km	7.1.3	AA1403011-E6

Transceivers		Minimum software version	Part number
10GBASE-ER	1550 nm SMF with a range up to 40 km	7.1.3	AA1403013-E6
10GBASE-SR	850 nm with a range up to: <ul style="list-style-type: none"> • 22 m using 62.5 micrometer (μm), 160 megaHertz times km (MHz-km) MMF • 33 m using 62.5 μm, 200 MHz-km MMF • 66 m using 62.5 μm, 500 MHz-km MMF • 82 m using 50 μm, 500 MHz-km MMF • 300 m using 50 μm, 2000 MHz-km MMF • 400 m using OM4 MMF 	7.1.3	AA1403015-E6
10GBASE ZR/ZW SFP+	1550 nm 70km SMF. The achievable link distance is primarily dependent on cable plant insertion loss. 70 km is not possible in some situations.	7.2.10	AA1403016-E6
10GBASE-LRM	1310 nm. Up to 220 m reach over Fiber Distributed Data Interface (FDDI)-grade 62.5 μm multimode fiber. Suited for campus LANs.	7.1.3	AA1403017-E6
10GBASE-CX	4-pair direct attach twinaxial copper cable to connect 10 Gb ports. The maximum range is 10 m.	7.1.3	AA1403018-E6
10GBASE-CX	4-pair direct attach twinaxial copper cable to connect 10 Gb ports. The maximum range is 3 m.	7.1.3	AA1403019-E6
10GBASE-CX	4-pair direct attach twinaxial copper cable to connect 10 Gb ports. The maximum range is 5 m. t —	7.1.3	AA1403020-E6
10GBASE-ER CWDM SFP+	1470 to 1610 nm with a range up to 40 km	7.2.10	AA1403153-E6 to AA1403160-E6

Table 4: Module and transceiver compatibility

Module	100BASE-FX SFP	1000BASE-XX SFP	10GBASE-XX XFP	10GBASE-XX SFP+*
8812XL	—	—	—	✓
8834XG	✓	✓	✓	—
8848GB	✓	✓	—	—
8848GT	—	—	—	—
8612XLRS	—	—	✓	—
8634XGRS	✓	✓	✓	—
8648GBRS	✓	✓	—	—
8648GTRS	—	—	—	—
8630GBR	—	✓	—	—
8648GTR	—	—	—	—
8683XLR	—	—	✓	—
8683XZR	—	—	✓	—
* Includes 10GBASE-CX direct attach copper cables				

Unsupported hardware for Release 7.2

Release 7.2 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 7.2 supports the 8692 SF/CPU only if it is equipped with SuperMezz. The 8692 SF/CPU without SuperMezz is not supported with Release 7.2.

In addition, M mode is no longer supported in Release 7.2. The software runs in R mode by default.

 **Important:**

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 7.2. 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, 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 similar in behavior to the 802.3ad standard for static LACP.

*** Note:**

For the latest information on the Unified Access scaling capabilities, refer to the WLAN 8100 release notes.

Table 5: Supported scaling capabilities

	Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU
Shortest Path Bridging MAC (SPBM)	
ARP entries with SPBM enabled per switch (tested in ERS Release 7.1.0.0)	6000
C-VLANs to I-SIDs supported with SPBM enabled per switch without SMLT	2000
C-VLANs to I-SIDs supported with SPBM enabled per switch with SMLT	1700
VRF instances	256 (including GRT)
GRT routes with SPBM enabled per switch (tested in ERS Release 7.1.0.0)	8000
VRF routes with SPBM enabled per switch (tested in ERS Release 7.1.0.0)	8000
L2 VPNs	2000
Number of MACs on VPNs	30 000
Layer 2	
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

	Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU
SMLT IDs	127
SLT IDs	382
VLANs on SMLT/IST link	With Max VLAN feature enabled: 2000
RSMLT links per VLAN	64
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
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 brouter-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)	1024
Static route entries	2000 per VRF (10 000 per system)

	Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU
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 per VRF (200 per system)
OSPF routes	20 000 per VRF (50 000 per system)
OSPF interfaces	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>	
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 without SMLT	4000 (per switch)
Multicast streams with SMLT	3000 (per switch)
Multicast streams per port	1000
Multicast streams on non-SPBM VLANs when SPBM is enabled on the switch	1500
IGMP reports/sec	250

	Maximum number supported using 8692 SF/CPU with SuperMezz or 8895 SF/CPU
IPv6	
IPv6 interfaces	250
IPv6 tunnels	350
IPv6 static routes	2000
OSPFv3 areas	5
OSPFv3 adjacencies	80
OSPFv3 routes	5000
Operations, Administration, and Maintenance	
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
RMON Ethernet statistics with 4000K memory	4590
BFD	
BFD sessions per switch	256 (This is a combined total of IPv4, IPv6, and VLACP sessions.)
LLDP	
LLDP neighbors per port	16
LLDP neighbors per switch	12288

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 6: License and features

Base License	Advanced License	Premier License
<ul style="list-style-type: none"> • Avaya VENA Unified Access • IP Multinetting • IP Source Guard • DHCP Snooping • Dynamic ARP Inspection • BPDU Filtering • IGMP Querier for L2 • PIM-SSM for SMLT • Multicast VLAN Registration (MVR) • LLDP 	<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"> - IPv6 Routing - IPv6 over SMLT and RSMLT - DHCPv6 Relay - VRRPv3 - BGP+ - RADIUSv6 - BFD over IPv6 	<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 IP Shortcuts (VRF0 shortcuts) - SPB L3 VSNs (VRF Extensions) - IP VPN Lite over SPB IP shortcuts - Inter-VSN Routing - IEEE 802.1ag Connectivity Fault Management - IP Multicast over SPBM

All IPv6 features require the Advanced License.

Ethernet Routing Switch 8800/8600 Release 7.2 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).

Important operational notes

Keep the following notes in mind when operating the ERS 8800/8600:

- Do not use TFTP to transfer large files to or from the device. There is a possibility of causing stack overflow and potentially crashing the system. Use FTP instead.
- Keep the out-of-band management network addressing scheme completely separate from the in-band network addressing. This separation ensures that no overlap results. Any overlap between the management network and the in-band network can result in network loops and severely affect network performance.

Avaya does not support configurations with any overlap between the management and user address schemes.

- Use the CLI/ACLI **verbose** mode feature to generate a verbose mode configuration for the sole purpose of viewing default configuration options.

Avaya does not support sourcing or loading verbose configuration files. Use non-verbose configuration files only for device operation.

- Do not configure non-IST inter-switch links running parallel to an inter-switch IST. This can result in network loops and is not a valid configuration. You should never use this kind of configuration.
- When upgrading your system, make sure the image files match.

System upgrades that result in a condition where the system boots with mismatched boot-monitor image (“b” or “be” image) and run-time image (“a”, “ae” or “m” image) can result in incorrect system behavior. It is very important to follow the correct procedure to boot with the boot-monitor image.

For more upgrade information, see *Avaya Ethernet Routing Switch 8800/8600 Upgrades* (NN46205-400).

- Take special care when upgrading dual CP systems running in hot-standby HA mode. The two CPs must be booted in a specific order and within a very specific time gap between them. Booting outside of these guidelines can result in upgrade failure. Furthermore, the upgrade of I/O cards will not take place if this procedure is not followed correctly. This results in a software release mismatch between the CP and I/O cards, which results in malfunction of the system.

For more upgrade information, see *Avaya Ethernet Routing Switch 8800/8600 Upgrades* (NN46205-400).

- Show commands such as `show sys info` that access system level information or SNMP gets/walks that access the system MIB should never be run from or against the Backup CPU.
- From the time the link establishment process begins, the 10G links on 8812XL modules may take a minute or so to become operational.
- SMLT/IST ports enter a special “hold-down” state transition immediately after reboot as the system comes up and until it reaches full operational state. With this in mind, do not make any dynamic configuration changes or movement of SMLT/IST ports until the system is completely up and all ports have reached their normal runtime states.

File names for this release

This section describes the Ethernet Routing Switch 8800/8600 Software Release 7.2.10 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 7.2, NN46205-400*.

Table 7: Release 7.2.10 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_72101.tar.gz	63,870,553
Copyright file	Ethernet Routing Switch 8600/8800 Master Copyright file	ERS8k.7.2.10.1_Copyright.docx	185,956
Ethernet Routing Switch images			
Boot monitor image for 8692 SF/CPU	8692 CPU and switch fabric firmware	p80b72101.img	1,197,268
Boot monitor image for 8895 SF/CPU	8895 CPU and switch fabric firmware	p80be72101.img	1,267,788
Run-time image for 8692 SF/CPU	Run-time image for 8692 SF/CPU	p80a72101.img	17,009,428

Module or file type	Description	File name	Size in bytes
Run-time image for 8895 SF/CPU	Run-time image for 8895 SF/CPU	p80ae72101.img	15,945,263
Run-time image for R modules	Image for R modules	p80j72101.dld	1,909,176
Run-time image for RS and 8800 modules	Run-time image for RS and 8800 modules	p80k72101.dld	1,977,308
Run-time image for Enterprise Enhanced SF/CPU Daughter Card (SuperMezz)	Image for the SuperMezz card	p80m72101.img	17,086,410
3DES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c72101.des	56,124
3DES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce72101.des	51,972
AES for 8692 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80c72101.aes (this image includes the 3DES image)	27,644
AES for 8895 SF/CPU	Encryption module for privacy protocol with Secure Shell (SSH)	p80ce72101.aes (this image includes the 3DES image)	25,316
MIB	MIB files	p80a72101.mib	5,678,451
MIB (zip file)	Zip file containing MIBs	p80a72101.mib.zip	896,217
MD5 checksum file	md5 checksums of all Release 7.2 software files	p80a72101.md5	1,195
Firmware images			
FOQ for R modules	Feedback output queueing FPGA firmware	foq267.xsvf	5,320,469
BMC for R modules	BAP memory controller FPGA firmware	bmc776.xsvf	2,640,266

Module or file type	Description	File name	Size in bytes
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 and 8800 modules	Contains FOQ, BMC, DPC, mirroring, and loopback images	rs_dpm_fpga.bin	4,538,368
PIM images for RS and 8800 modules*	PIM FPGA firmware required for the 8612XLRS module	pim8612XLRS.bin	60,183
	PIM FPGA firmware required for the 8634XGRS and the 8834XG modules	pim8634XGRS.bin	78,173
	PIM FPGA firmware required for the 8648GBRS and the 8848GB modules	pim8648GBRS.bin	79,891
	PIM FPGA firmware required for the 8648GTRS and the 8848GT modules	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	ers8000v7210_HELP_EDM_gzip.zip	3,995,688
ERS 8000/8600 EDM plug-in for COM			
EDM plug-in for COM	EDM plug-in for COM	ers8000v7.2.10.0.zip	5,343,483

* With the `pim8612XLRS.bin` file, you can upgrade older versions of the 8612XLRS module to PIM (1024). However, there is no PIM upgrade file available for the 8812XL. It is programmed at the factory at 256.

Table 8: Firmware versions for modules

Module Type	CC	FOQ	DPC	BMC	PIM	Mirror	Loopback
8683XLR and 8683XZR	3	267	194	776	0*	N/A	N/A
8630GBR	3	267	194	776	769	N/A	N/A
8648GTR	3	267	194	776	3*	N/A	N/A
8812XL	33281	270	7	264	256	304	274
8612XLRS	33281	270	7	264	1024	304	274
8648GBRS and 8848GB	33281	270	7	264	1024	304	274
8648GTRS and 8848GT	33281	270	7	264	768	304	274
8634XGRS and 8834XG	33281	270	7	264	1024	304	274

* No PIM image files available

Important information and restrictions

This section contains important information and restrictions that you should consider before you upgrade to Release 7.2.

Fixes from previous releases

The Ethernet Routing Switch 8800/8600 Software Release 7.2 incorporates all fixes from prior releases up to and including release 5.x, 7.1.x, and 7.2.1x.

8812XL SFP+ I/O module upgrade considerations

- When you upgrade other modules, the 8812XL SFP+ I/O module displays KHI errors.

! Important:

Avaya recommends taking the 8812XL SFP+ I/O module out of service before you upgrade other modules.

- There is no PIM upgrade file available for the 8812XL. It is programmed at the factory at 256.

SuperMezz, SF/CPU memory, and upgrades

To support Release 7.2, the 8692 SF/CPU must be equipped with SuperMezz. 8692 SF/CPU without SuperMezz is not supported with Release 7.2. If the Release 7.2 software is booted with a non-SuperMezz 8692 SF/CPU, the line cards do not come online.

For Release 7.2, 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 ensure proper software cleanup on the CP and to prevent corruption of the external compact flash card or the PCMCIA card, **do not remove the external memory card without first entering the following command:**

- `dos-stop /pcmcia`

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-chkdisk <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, and the boot monitor.

 **Note:**

When the number of files stored on PCMCIA card or internal flash exceeds 100, the I/O modules may experience a delay in coming online only during an upgrade or downgrade. If there are more than 100 files on the PCMCIA card or internal flash, delete any unnecessary files.

Proper handling of SF/CPU and I/O modules

 **Caution:**

Avaya strongly recommends that you disable any module (SF/CPU or I/O) before you remove it. Use one of the following commands.

- `config slot <slotnum> state disable` (CLI command)
- `slot shutdown <slotnum>` (ACLI command)

Do not remove any module without first entering one of the preceding commands.

For more information on using these commands for specific tasks, see the following topics in *Upgrades* (NN46205–400):

- Upgrading from 8692 SF/CPU with SuperMezz to 8895 SF/CPU
- Hot swapping the Master SF/CPU module in a dual CPU chassis
- Hot swapping the Secondary SF/CPU module in a dual CPU chassis
- Hot swapping an I/O module

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, versions 8.x and 9.x
- Microsoft Internet Explorer, versions 8.x and 9.x

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. The EDM Web server is disabled by default and must be enabled through CLI or ACLI.

Important:

Because the EDM Web server is disabled by default for security consideration, you must enable it and change its default user password with CLI/ACLI commands. For more information, see *Avaya Ethernet Routing Switch 8800/8600 — User Interface Fundamentals* (NN46205-308).

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.

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 replaces older graphical user interfaces

Enterprise Device Manager (EDM) replaces both the Java-based Device Manager and Web-based management. EDM is an embedded element management and configuration application for Ethernet Routing Switch 8800 Series switches. EDM uses a Web-based graphical user interface for the convenience of full integration onto the switch, but it retains the look and feel of Device Manager.

! Important:

With the introduction of Enterprise Device Manager (EDM), the use of Device Manager (sometimes referred to as JDM) is no longer supported because the use of JDM to control the switch could lead to potential corruption of the switch configuration.

! Important:

If you upgrade the software on your switch, and if you are managing the switch with EDM, then you should refresh the browser cache on your end device to ensure that EDM loads the latest tabs for all respective features.

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 7.2, Avaya recommends using COM 2.3, 3.0 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 7.2 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_71_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.

DiffServ consideration with IPv6 inter-tunnel routing

ERS 8800 7.2.10 release supports inter-tunnel routing of IPv6 packets encapsulated in an IPv4 tunnel. Packets routing **from one tunnel to another tunnel** do not comply with the access DiffServ policy, which always expects the incoming port to be an L3-trusted interface.

This limitation applies only to incoming tunneled packets that are expected to be routed into another tunnel.

HSECURE considerations

When the `hsecure` feature is enabled, switch access is limited to SSH only. All other remote access methods including SNMP are disabled.

An alternate method to limit access to the switch is to use boot flags to individually disable undesired access methods such as `telnet`, `ftp`, and `rlogin`. To block SNMP, enable the `block-snmp` boot flag. To limit SNMP access to SNMPv3 users, the switch administrator must delete the default community strings.

Supported upgrade paths

The Ethernet Routing Switch 8800/8600 Software Release 7.2 supports direct upgrades from the following earlier releases:

- 5.1.1.10
- 5.1.1.11
- 5.1.8.x
- 7.0.0.3
- 7.1.x.0
- 7.2.0.0
- 7.2.0.2

If you want to upgrade to release 7.2.10 from any other release, first upgrade to one of the above releases and then upgrade to 7.2.10.

General upgrade considerations

The configuration file generated with Ethernet Routing Switch 8800/8600 Software Release 7.2 contains options that are not backward-compatible with any previous Ethernet Routing Switch 8800/8600 Software Releases.

Loading a Release 7.2 configuration file on a pre-7.2 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 7.2

Before you upgrade, read *Avaya Ethernet Routing Switch 8800/8600 Upgrades, 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 42.

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

Upgrading R modules from Release 5.1.1.10/11 to 7.2.10.1

Use this procedure to upgrade the R modules in your switch.

Important:

This procedure is for R modules **only**. If you have any RS modules that require upgrading, use the following procedure but use the RS `bin` modules provided via software.

Procedure

1. Use the following command to verify that the `mezz` flag is set to `true`:

```
show boot config flags
```

2. Check the version of the Dual Port Controller (DPC) FPGA firmware for the 8630GBR, 8648GTR and 8683XLR/8683XZR.

```
show sys info asic
```

Note:

If the version is below 184, follow Step 3.

3. Upgrade the FPGA firmware. During this process, do not press `Enter` or interrupt this process.

Note:

In the following example, `<slot-number>` refers to the slot that the module is in so enter the appropriate number.

```
config diag update fpga dpc <slot-number> /flash/dpc184.xsvf
```

Make sure you receive a log message on the console stating that the upgrade of the FPGA was successful.

4. Repeat Step 3 for any other modules that need upgrading of their DPC FPGA.
5. Repeat Steps 2 and 3, but replace DPC with the appropriate FPGA type to upgrade the FOQ, BMC, and PIM firmware.

Note:

PIM is not applicable for the 8648GTR.

6. Set the `mezz` flag to `true`:

```
config boot config flag mezz true
```

7. Save the boot configuration file.
8. Boot the switch and press `y` to confirm.

```
boot
```

9. After the switch restarts, verify that the ASIC is the correct version make sure the SuperMezz is detected and it is up.

```
show sys info asic
```

```
show sys info
```

10. Copy the 7.2.10.1 images to flash and set the primary image file to point to p80a72101.img.
11. Save the boot configuration file.
12. Restart the switch.

*** Note:**

Do not interrupt the boot ROM upgrade process.

```
boot /flash/p80b72101.img
```

13. **After the boot rom is upgraded and the system is ready to boot /flash/p80a72101.img, interrupt the process if the /flash was formatted to a pre 5.0.0.0 version and follow these steps:**
 - a. Copy the necessary files over to pcmcia.
 - b. Format the flash.
 - c. Copy all the files back to /flash.
14. At the monitor prompt, enter the following command and let the switch boot through normally.

```
/flash/p80a72101.img
```
15. After the system is fully functional in 7.2.10.1, repeat Step 3 at the console prompt to upgrade the line cards to dpc version 194. However, instead of dpc 184, you are going upgrade to 194. During this process, do not interrupt the upgrade. After you see a message stating that the upgrade of the FPGA was successful, repeat this same procedure for the rest of the modules.
16. Reboot the switch and press `y` to confirm.

```
boot
```

17. After the system is fully operational, verify that all the R modules dpc is set to 194.

```
show sys info asic
```

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 7.2 code with the 8692 SF/CPU with SuperMezz as master and slave.
- Download the p80ae7200.img and p80be7200.img software images, as well as the dld files (p80j7200.dld, p80k7200.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 (p80ae7200.img, p80be7200.img, p80j7200.dld, p80k7200.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 p80ae7200.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/ p80be7200.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 (p80ae7200.img, p80be7200.img, p80j7200.dld, p80k7200.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/ p80be7200.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 7.2, 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* (NN46205-400).

Upgrade considerations: Power Management

The Power Management feature available with Release 7.2 may require you to take special steps before you upgrade.

When you upgrade to Release 7.2, Power Management is enabled by default. When the switch boots, 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.

 **Note:**

Avaya recommends using the power supply calculator to determine if the 8005AC/8005DC (Single or Dual Input) power supplies are required. The 8004AC power supply can be used with R modules and is supported in Release 7.2.

 **Important:**

The 8004AC power supply runs the PSUs @ 110VAC/15A. When you upgrade from the 8004AC and/or DC power supplies to the 8005AC and/or DC power supplies, be aware that the recommended input voltage is 200-240VAC to obtain full output power from 8005 power supplies. Additionally, 20AMP circuits @ 110V are required. Therefore, review or update your Power Plants and UPS accordingly.

For Power Management configuration and conceptual information, see *Avaya Ethernet Routing Switch 8800/8600 Administration, NN46205-605*.

Power Management operations

When the switch boots with Power Management, 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.

Upgrading systems with 8895 CPUs and Dual Input Power Supplies

In systems running pre-7.1.5.2 code with an 8895 SF/CPU and Dual Input Power Supplies, the I2C bus can sometimes hang. The system continues to function properly even with the I2C in a hung state. However, when you reboot the system, the reboot fails due to `Chassis Seeprom read error`.

Important:

If the system that you are going to upgrade is already in the I2C Hung state, use the procedure below to correct that state **before** proceeding with the upgrade. Failure to do so results in a failed upgrade and the device remains in a software hung state upon reboot and potentially become inaccessible.

The I2C Hang issue is not a problem in the following code streams:

- Release 7.1.5.2 and higher in the 7.1.x software stream
- Release 7.2.1.1 and higher in the 7.2.x software stream

Note:

The Dual AC-Input Power Supply part number is: 8005DI 1170W/1492W 120/240V.

Before you begin

On the Primary CPU only, enter the CLI `show sys info` command (or ACLI `show sys-info`) to check the I2C status. If the Fan temperature reading is 0, that means the device is in an I2C Hung state and you must perform the following procedure. If it is not hung, you can skip this procedure and continue with the upgrade.

Procedure

1. Make sure someone is **physically** present at the device in case a power cycle becomes necessary to recover the system.

2. Download the 7.1.5.2 or higher software release onto the device and complete all the pre-sets as you would in a normal upgrade.

! Important:

DO NOT issue the `reset` or `reboot` command yet.

3. Once all pre-work is complete, power the chassis off.
4. Remove **ALL** power supplies from the chassis and keep them out for a minimum of **30** seconds.
This step discharges the capacitors in the supply, which makes the supply completely cold electrically. This step ensures that the I2C bus starts clean on power-up.
5. Reinsert all the necessary power supplies into the chassis securely.
6. Power on the chassis.
The chassis should boot successfully with the 7.1.5.2 or higher or 7.2.1.1 or higher software that it was preconfigured to load. If for any reason the chassis hangs during software initialization of this upgrade boot, you can recover it using the above procedure.

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 and 8800 I/O 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-7.2 or current 7.2 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
#          FQO=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 7.2 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 7.2.0.0. 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 7.2.0.0 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, 7.1.0.0 or 7.2.0.0 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)
```

High Availability mode considerations

Switches with two SF/CPU use High Availability (HA) mode to recover quickly if one SF/CPU fails. High Availability mode (also known as HA-CPU) permits the synchronization of configuration and protocol states between the Master and Secondary CPUs.

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

Important:

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

By default, HA-CPU is disabled in Release 7.2. To enable it enter the following command:

```
config bootconfig flags ha-cpu true
```

After you enable High Availability mode, the secondary SF/CPU resets to load settings from the saved boot configuration file. You must reset the primary SF/CPU after the secondary SF/CPU completes booting.

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

- PGM

PCMCIA considerations

Downloading files (40 MB or larger) to an ERS 8800/8600 from a PCMCIA card may impact the operation of your network.

Important:

Avaya recommends that you download large files to or from a PCMCIA card during off-hours or a maintenance window to minimize impact on the network.

Ongoing considerations

The following sections describe considerations that are not new for Release 7.2, but which still apply for 7.2.

Module and chassis compatibility and performance considerations

Release 7.2 does not support classic modules. Only R, RS, and 8800 series line card modules are supported with release 7.2. Also, the 8003 chassis is not supported with release 7.2. 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 7.2.

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 9: 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
8812XL	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 and RS 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 10: Chassis hardware revision

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

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 11: Switch behavior during boot cycle

Parameters	Switch action
A configuration file is not specified. The <code>config.cfg</code> file is present on the flash drive.	The switch boots <code>config.cfg</code>

Parameters	Switch action
The primary configuration file is specified. The configuration file is present on the flash drive.	The switch boots the specified configuration file.
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 image-name /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

Consider the following issues when configuring IPv6.

Learning addresses on multiple interfaces

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:
```

Triangle SMLT

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.

Inter-tunnel routing

Previous releases of the ERS 8800/8600 supported IPv6 traffic over IPv4. Release 7.2.10 expands on this capability by supporting tunnel-to-tunnel traffic, which is sometimes called *6in4 tunneled traffic*.

The ERS 8800/8600 now supports inter-tunnel routing of IPv6 packets encapsulated in an IPv4 tunnel. However, the packets routing from one tunnel to another tunnel do not comply with the access DiffServ policy, which always expects the incoming port to be an L3-trusted interface. This limitation applies only to incoming tunneled packets that are expected to be routed into another tunnel.

*** Note:**

The ERS 8800/8600 switch also supports the following cases:

- You can provision multiple tunnels on the same switch where traffic is de-capped and natively routed in the IPv6 network.
- The same ERS 8800/8600 switch can serve as the starting point for multiple 6in4 tunnels.

BFD over RSMLT

BFD over RSMLT is not supported.

Deletion of BFD sessions

To delete a Bidirectional Forwarding Detection (BFD) session, disassociate all the applications with the BFD session and then administratively bring down the BFD session.

To successfully delete a BFD session, you must execute the commands in the following order:

1. Disassociate all applications from the BFD session.
2. Disable BFD at the global or interface level, which transitions the BFD session to the `ADMINDOWN` state.

*** Note:**

If you change the above order of operations, it does not delete the BFD session.

For more information, see the Bidirectional Forwarding Detection chapter in either the IPv4 manual (*Avaya Ethernet Routing Switch 8800/8600 Configuration — IPv4 Routing NN46205–523*) or the IPv6 manual (*Avaya Ethernet Routing Switch 8800/8600 Configuration — IPv6 Routing NN46205–504*).

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 7.2 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 soon). Otherwise, traffic loss can occur.

SMLT considerations

Software Release 7.2 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, 7.1.0.0, and 7.2.0.0.

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

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.

 **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 5.1.x release are required, versus a direct upgrade to 7.2.0.0. 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 7.2.0.0 are supported from 4.1.8.x, 5.0.1, 5.1.2, 5.1.3, 5.1.4, 5.1.5, 7.0.0.0, 7.1.0.0, 7.1.1 and 7.1.3.0.

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 7.2.0.0.




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 7.2.0.0. Due to the SMLT behavior, the network diagram now looks like the following figure.




Figure 2: SMLT upgrade scenario after upgrading switch A to 7.2.0.0

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 7.2.0.0, 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 7.2.0.0 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 4.x, or 5.0.0.x (where x is less than 2), the switches **MUST** be upgraded to 5.1.x before upgrading to 7.2.0.0.

Modifying IST ports during an SMLT upgrade

! Important:

Do not modify the IST ports immediately after rebooting the SMLT-enabled node.

When an SMLT enabled device is rebooted, SMLT/SLT and all normal ports enter a special "hold-down" state transition and remain locked till the system becomes fully operational. Addition of any port that is in such an "hold-down" state to an IST will cause continuous flooding of MLT unlock messages such as `MLT INFO smltTick: Initial MAC/ARP tbl completed, unlocked SMLT/SLT ports` to the log.

To avoid this scenario, any configuration changes on or modification of IST ports should be done only when the system is completely up, IST is up and all ports have reached their normal run-time states. Even if IST fails to come up all ports will reach their normal run-time state in 40 seconds after the system is completely up.

*** Note:**

Avaya recommends using the following procedure to modify IST ports during a node upgrade.

Procedure

1. Keep the desired new IST ports in an administratively down state and ensure connectivity.
2. Disable the IST using the following CLI command:
`config mlt <mltID> ist disable`
3. Disable all of the other ports on the node.
4. Remove the old ports from the IST MLT:
`config mlt <mltID> remove ports <ports>`
5. Add the new ports to the IST MLT:
`config mlt <mltID> add ports <ports>`
6. Enable the new IST ports:
`config ethernet <slot/port> state enable`
7. Verify that the IST is up:
`show mlt ist info`

8. Enable all other required ports on the node.
 9. Check all the SMLT/SLT's on the node.
 10. Save the configuration.
 11. Upgrade the node using the new configuration file.
-

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.

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 failedfor 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 !!
```

Avaya recommends using the Enterprise Policy Manager to simplify operations with the centralized management of ACLs and ACEs.

MTBF for 1 Gig SFPs and 10 Gig XFPs

The mean time between failure (MTBF) for all 1 Gig SFPs is 807 000 hours. The MTBF for all 10 Gig XFPs is 675 000 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 7.2, 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 7.2, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting*, NN46205-703.


Chapter 4: Resolved issues in 7.2.10

This section details all the issues that were found in previous releases but resolved in Release 7.2.10.


Table 12: Resolved issues in 7.2.10


WI reference	Description
wi00508462	In some cases, the output of the <code>show slpp interface gig</code> command does not show anything on either IST peer even when SLPP has brought the ports down. The command should normally display some information on either one or both of the IST peers.
wi00955243	If VLACP is deployed over ISIS/SPB interfaces and the ISIS/SPB interfaces are MLT interfaces, when an HA-mode CPU switchover is performed VLACP brings down all the ISIS MLT interfaces and hence the ISIS adjacencies with them.
wi00968354	L2 Ping for CFM - CMAC does not work when all the links connecting to one of the SPBM core switches are disabled. Workaround: Disable Common and Internal Spanning Tree (CIST) on the NNI ports.
wi01003993	Fixed an issue with IPFIX where the data would be dropped if the active timeout and the expiry intervals happened to be exact multiples of each other.
wi01026921	After disabling LACP on an interface, the SLPP PDUs are not getting bridged .
wi01033794	In the CLI, you can make a stub area and then remove it by disabling the stub on the area. In the ACLI, there is no way to remove a stub area. Workaround: Use the default area x.x.x.x stub, which disables the OSPF stub similar to the CLI disable.
wi01033833	Rebooting the switch or resetting the 8812 line card causes link flap and automatic port shutdown. Workarounds — There are two workaround options: <ul style="list-style-type: none">• Disable link flap detect auto port down (enabled by default) (not recommended)• Increase the link flap detect interval to a number greater than the duration of the link flap seen.
wi01035348	Ping and traceroute are not functional when sourced from one VRF to a device of another VRF.
wi01035373	The ERS 8800 is unable to apply a route policy to the redistributed routes from one VRF to another. Setting the redistribute table route map fixed this issue.
wi01035599	For a dual-stack-enabled (IPV4 and IPV6) VLAN with RSMLT edge support enabled, upon an HA failover, the RSMLT timers get reset and start from initial value instead of continuing from current value. This issue has been fixed. Also fixed

WI reference	Description
	the problem where RSMLT holdup timer not counting in edge support scenario and peer MAC is made local permanently.
wi01036125	Column formatting is skewed for the output of <code>show ip route vrf 1 spbm-nh-as-mac</code>
wi01036150	Column formatting is skewed for the output of <code>show ip igmp interface</code>
wi01036998	Unable to remove PCAP filters as license no longer valid, preventing port mirroring functionality.
wi01037161	When you try to delete a VRF ARP entry in GRT mode, an invalid error message displays (CPU6 [08/14/12 06:27:01] IP ERROR VRF name: vrf10 (VRF id 10): rcIpDeleteStaticArp:unable to find Static Arp R ec for 3.5.2.220IfIndex 2400) even though the ARP entry exists and the ARP entry is deleted successfully.
wi01037846	In-service replacement of an IST peer chassis could result in a condition where the non-replaced chassis continues to use the wrong (older chassis) MAC address in the SLPP packets. This issue has been fixed. Also fixed a potential bad memory access situation in the MAC delete by port function.
wi01038103	The CLI and ACLI no longer support the <code>min-lsp-gen-interval</code> command. In EDM, you can still set this command (<code>MinLSPGenInt</code>) from the IS-IS > IS-IS > System Level tab, but the switch returns a <code>Not Supported</code> message.
wi01038842	Syslog is producing incorrect Host Numbers when enabling or disabling the service. Workaround: Use the corresponding CLI command <code>config sys syslog host <id> host disable</code> .
wi01041491	<code>show fulltech</code> can be disallowed if SSH session is stranded during execution.
wi01041501	ERS 8800 Radius CLI Profile now works with SSH.
wi01041506	Inconsistent ACL/ACE behavior is observed on 8648GTRS module.
wi01041511	ERS 8800 with CP card 8895SF/CPU registered three consecutive core dumps. Device running on s/w 7.1.3.0. SOP init problem.
wi01041516	ACL problem: NextHopReachableRedirect - When the redirection destination is not reachable, the filter continues to redirect instead of dropping.
wi01041521	Syslog message was sent using LAST created CLIP Interface IP address instead of the first configured CLIP address.
wi01041523	ACLI: Access-policy-SNMPv3 configuration disappears after reboot.
wi01041685	In Unified Access Networks with multicast traffic flowing from wireless devices, the part of the VLAN using PIM-SSM and IGMP version 3 does not reset the <code>copy to cp</code> flag when it sends the first packet to the CP. Hence, all the packets from those flows are sent to the CP resulting in high CP utilization that may affect other services

WI reference	Description
	<p> Note: This issue can cause high CP utilization when traffic is flowing more than 50 Mbps from the wireless IGMPv3 senders.</p>
wi01041714	The 8895 SF/CP sends out BFD control packets too early.
wi01042922	A possible crash could happen when running <code>show fulltech</code> during <code>show ip count-summary</code> section.
wi01044664	Device sends LACP with irregular timer on average every 15 sec when configured for long timer of 30 seconds. This was caused by two independent tasks generating the LACP PDUs. Added synchronization between these tasks to ensure a task does not send the PDU if other task has sent it for the specified interval.
wi01046166	Upon user-initiated device reset from CLI, the logging operation waiting to flush the logs before proceeding to reset can cause a deadlock and force a watchdog reset.
wi01046772	SMLT/LACP sends aa:aa:aa:aa:aa:xx for Actor System Id.
wi01051880	Traffic does not pass between a server with Intel 10G NIC connected to an 8812 after the server is rebooted.
wi01054153	Routed packets that should be routed out the WAN VRF (Default Gateway) are routed back and forth between the two Office ERS 8600s.
wi01054991	Security Scan (Nessus) reports Telnet and FTP ports open although service disabled.
wi01055421	When an IST link is physically removed, the IST does not recover if the link is left down for more than 8 minutes.
wi01058481	Removing and reinserting a line card causes inter-VRF routing issues.
wi01058539	The nick-name option in <code>show isis spbm multicast-fib</code> sometimes displays the system name for the wrong BEB. A code change to compare three bytes of the nick-name instead of two fixes this issue.
wi01058709	With DNS Setting configured, running backup scripts via Telnet does not execute all the commands.
wi01061665	Under certain conditions, an ERS 8600/8800 device that has IPFIX enabled could see the Line Card (COP) CPU utilization spike to 100%, affecting other protocol operations or resulting in network instability/slowness. That issue has been fixed. An issue related to the accuracy of the data reported by IPFIX has also been corrected.
wi01064178	When an ARP request is sent out for the NLB virtual IP address (VIP) in the "nlb multicast" mode, all the hosts within that cluster reply back with a multicast MAC address. The ERS 8600 was adding the MLT/port on which the reply is received. On one of the SMLT peers, this is going to be the IST. Henceforth, any traffic destined to the VIP from that ERS will be sent on the IST and that will be dropped


WI reference	Description
	by the peer. To fix this issue and prevent the peer from dropping traffic, the ERS now learns the port/MLT of the SRC Mac of the host sending back the ARP reply.
wi01064252	Under high levels of multicast traffic, the 8630GBR experiences resets due to false packet memory refresh errors.
wi01065076	COP Record corruption and subsequent crash could result due to a race condition in the handling of FTN and IP routes in the RTM.
wi01065450	The CP-Limit feature does not work on ports on slots 8 through 10 on release 7.1.x.x (older than 7.1.5.1) and release 7.2.0.2. Fixed the CP-limit feature to shut the offending ports down for slots 8, 9 and 10.
wi01065671	An HTTPS-related crash can result from running out of file descriptors.
wi01065676	A null pointer crash in the BGP area resulting from processing of a state change event has been fixed.
wi01066167	Port does not initialize correctly with a 100Base FX SFP and the following message is reported: KHI Warning Port <number> is experiencing Packet Errors, FCS Errors.
wi01066759	Crash seen on ERS 8800/8600 when performing VPFM discoveries against ERS 8800/8600 seeds has been fixed. This crash can occur when the ERS 8800/8600 has both AES and 3DES encryption modules loaded and VPFM includes both SSH and SNMP credential configuration.
wi01071995	A user with read-only privilege is not able to launch a non-Global VRF Context View using the COM and EDM Plug-in. EDM Plug-in v7.2.1.1 addresses the above issue. A read-only user can now launch a non-Global VRF Context View from the COM Application.
wi01076837	In some cases an incorrect MAC address value is displayed in the SLPP log entries related to loop detection.
wi01073250	Removing and re-adding a 10GB Link to the STG is not recalculating the physical path cost to the expected value of 1.
wi01075445	The ACLI does not have a command for removing an OSFP stub area like there is in the CLI. This was resolved by adding a stub option under the NO form of the area command.
wi01075452	The VRF IP static multicast MAC ARP entries were causing two issues: an invalid error message and incorrect behavior. The invalid message appeared because a deleted entry was still appearing in the show command. A comparison check was added to avoid this issue. The incorrect behavior was that multiple exit entries were appearing in the configuration file. After an exit, the flag was set properly to prevent this from happening.
wi01077334	Traceroute is not working for IPv6 routes when a tunnel is configured in the network. IPv6 packets UDP checksum gets corrupted going onto IPv6 tunnel.
wi01077643 wi01018835	BGP message exchange between an ERS 8800/8600 peering with a host running Linux OS can result in a crash of the switch. Some Linux implementations are

WI reference	Description
	incorrectly sending update message with the “AS path attribute length extension” bit set even when the length is less than or equal to 255 bytes. The ERS 8800/8600 now handles this condition by adjusting the bit, if necessary, by checking the actual length of the update message.
wi01077643 wi01045784 wi01049417	ERS 8800/8600 switch can experience a CPU switchover due to memory corruption. Robustness fixes have been made in the areas of BGP, RADIUS, Route Redistribution, Secure Copy (scp) and Ethernet.
wi01077643 wi01060672	IGMPv3–related memory corruption has been fixed.
wi01077643 wi01062810	IPFIX configuration is only allowed in the default VRF context. IPFIX packets from certain slots could be wrongly interpreted as coming from/destined to a wrong VRF resulting in IPFIX packets not being sent out from certain slots. This issue is fixed by forcing all IPFIX collection packets to VRF 0.
wi01078774	Qualys scan related robustness fixes have been made.
wi01079116 wi01059293 wi01044656	After an OSPF FULL adjacency is established, if the last entry in the Link State Database has a nonzero forwarding entry, then any new external routes received in an LSUPDATE will be injected in to the routing table with that forwarding entry as their next hop.
wi01079122	When the switch receives a BGP Update message whose path attribute is greater than 260 bytes, it corrupts the memory area beyond 260 bytes and causes a crash.
wi01079127 wi01054153	After a reboot, the routed packets that should have routed out the VRF (Default Gateway) were routed back and forth between the two office ERS 8600s. This issue was fixed by adding an IP default route and checking the source VRF id.
wi01079131	If you insert a line card and the default route is on VRF 1, the following messages appear: COP-SW ERROR rcdAddEcmpDefaultRoute: Arp rcdRadixLookup failed and ercdProcArpRecMsg: Failed to Add Ecmp Default Route. A new function fixes this issue by installing the ARP records correctly.
wi01079445 wi01058836	<p>The error message Stop polling DI Power Supply due to excessive i2c error; fan temp=0 would be seen in system running 8895 and DI Power supplies. In systems with 8895 CPUs, there is a potential for the I2C bus to stall. This could result in peripheral devices on the I2C bus becoming unreachable from the CPU.</p> <p> Note:</p> <p>Take special care when upgrading a system running pre-7.2.1.1 software and having 8895s and Dual Input power supplies. For upgrade information, see Upgrading systems with 8895 CPUs and Dual Input Power Supplies on page 44.</p>

WI reference	Description
wi01078538	Intermittent BGP/VRRP/BFD session flaps when running CPU intensive commands such as <code>save config</code> and <code>show tech</code> with the <code>cli more</code> flag disabled.
wi01079634	When the last port of STG1 is removed, all ports added back to STG1 after config save.
wi01079637	Support for topology discovery for VSP4000.
wi01079638	Multiple PIM-SM crashes in memory instrumentation image.
wi01082834 wi01081346	A potential for system reset resulting from SNMP trapd task exception during a memory freeing operation has been corrected.
wi01083344 wi00853802	A metric set through the route policy was not applied to the IP shortcuts case. This affects both the GRT (TLV 135) and SPBM IP VPN (TLV 184). Added support to allow the metric to be manipulated through route policy. Now that you can configure the IS-IS TLVs, you can specify the same route from different edge BEBs advertising it.
wi01083344 wi01053163	IS-IS adjacencies up-time drifts by 5-6 seconds per minute.
wi01085280 wi01039790	The default trace level for the Infinity module was set to <code>VERY_TERSE</code> instead of <code>NO_DISPLAY</code> .
wi01085529	When you reset two BEBs that are IST peers, some of the L3 streams do not recover because the individual streams are using different ISIDs after the reboot. This issue was fixed by clearing the old LSP ISID information.
wi01085685 wi01046124 wi01066179	In ACLI mode, configuring an OSPF area range summary for the default area and saving the configuration results in an invalid saved configuration as the context was not set correctly to <code>router ospf</code> .  Important: Upgrade Warning: If upgrading a system that uses this configuration, the saved config from older release will need to be manually corrected and used for the upgrade to avoid configuration loss upon reboot.
wi01085689 wi01075385 wi00993971 wi01043317	In EDM, you cannot configure Inter-VRF route redistribution.
wi01085822	Enable/disable BFD at the interface level (VLAN, Port, and BFD interfaces) is hard coded in EDM.
wi01085998	Redistribution of a black hole static route results in a log error message: <code>SW ERROR Route 10.166.0.0/255.255.0.0 index 0 has null next hop pointer</code> . There was no impact to traffic flowing in the network. The problem was fixed by checking for black holes.
wi01086214	IPV6 admin state has to be toggled on vlan interface on static route vlan for the BFD session to be up.

WI reference	Description
wi01086324 wi01082670	When you create a black hole default route and reboot, it causes a system crash because RSMLT tries to add a temporary default route first. This issue was fixed by checking to see if the route being added is a black hole.
wi01086486 wi01086641	Running SPB while configuring IGMP Snoop on a VLAN caused a crash.
wi01086607 wi01086611	SSH/SSL-related memory leak and memory corruption has been fixed.
wi01087087	In EDM, an error message displays when you insert Static Route BFD NextHop.
wi01088133	When the MAC is learned on an NNI VLAN and that VLAN does not have an ISID configured on it, it was being incorrectly marked as NNI when the information was synced to the IST peer.
wi01088690	EDM display of the 8812XL module does not show in EDM after a reboot. SFP information does not show on reboot with administratively-disabled ports. SFP information was specifically bypassed for disabled ports to reduce polling overhead to ensure reads happen on all transitions.
wi01089721 wi01089705 wi01088321	The ARP delete operation has been fixed to ensure proper removal of ARP entries from hardware.
wi01089731 wi01085229	Security enhancement related to IPFIX feature has been added.
wi01090063	IPv6 Inter-Tunnel routing is not available on ERS 8600/8800.
wi01090869	The rcn[a]Sshxxx and rcn[a]AuthenticationSuccess SNMP Traps Deprecation are not supported in v7.2.10.0.
wi01091134	Deactivating and then activating a card causes an error message: COP-SW-ERROR logsCOP-SW-IP ERROR Slot 3: ercdProcArpRecMsg: Failed to Add ECMP Default Route .
wi01091954	Excessive flooding can result from an NNI outgoing interface being incorrectly programmed for NNI traffic.
wi01092905	IPv6 route failover issue where traffic fails every other time. The routes are present, but the traffic is not forwarded. IPv6 route does not revert after disabling/enabling a port.
wi01092917	No rcnAuthenticationSuccess SNMP trap defined in previous release MIBs.
wi01092921	Lack of rcLoginName in rcnaSshUnauthorizedAccess and rcnaSshSessionLogoutTrap name.
wi01092923	rcBFDTmpSessState in the rcnBfdStatusChange SNMP Trap is being received out of range.
wi01093784	rcnSshLogin Trap is NOT sent when an SSH session is set up BEFORE the wrong user credentials are entered.

WI reference	Description
	Proprietary rcnSshUnauthorizedAccess Trap is NOT sent for a wrong password input, instead, the standard SNMPv2 authenticationFailure Trap is sent.
wi01093789	IPv6 tunnels stop working after a period of time.
wi01094156	When the IS-IS system-id changes on the remote end without giving the local end a chance to declare the adjacency down, the IS-IS on the local node is kept in the INIT state and the remote node goes down. Because the local node still had the adjacency up when the hello packet from the remote came in with the new system-id, it was still sending the hello packets with the old system-id kept in its adjacency pointer, which was rejected by the other side. This issue was fixed by comparing the remote system-id in the adjacency pointer with the remote system-id in the header every time a hello message is received, and resetting the adjacency if it does not match
wi01094501	IGMP control traffic received on a UNI for a L2 VSN is sent to the CPU even though IGMP-SNOOPING or Multicast-Routing is NOT enabled on the VLAN. This behavior is incorrect. IGMP control traffic should only be sent to the CP if IGMP-SNOOPING or Multicast-Routing is enabled on the VLAN. This problem was a 7.2 issue only and it has been fixed.
wi01095181	The SW Datapath is being used for WLAN Tunnel packets on SPBM egress, not Unified Access.
wi01096332 wi01096320	Memory robustness improvements (fix memory leaks resulting from <code>show fulltech</code> command from performing SNMP operation on <code>rcPrFilterAceProtoShowTable</code> MIB have been fixed, from HA sync after creating filters for both IPv4 and IPv6 have been fixed).
wi01097042	Robustness improvement related to TACACS login.
wi01098085	When configuring RIP to redistribute IS-IS routes into RIP using a route-map, the configuration changes were not getting saved into the config file.
wi01098181 wi01095559	The <code>mlt id</code> argument on the <code>ip arp static-mcast</code> command does not load when loading the config after boot because it is not accepted by the syntax. This is specific to the VRF-based command. The GRT command works correctly. <pre>vlan static-mcastmac 4007 03:bf:0e:32:07:fc mlt 1,2,21,22 ip arp static-mcast 14.50.7.252 03:bf:0e:32:07:fc vid 4007 1,2,21,22</pre> The command is also putting 8 extra "exits" in the config file resulting in the rest of the config not getting loaded. This causes partial config loss on reboot.
wi01098738	Default route stopped working after OSPF adjacency flaps.
wi01098764	IPV6:RCIP6 ERROR rclp6StaticRouteDelRtm: failed deleting route from RTM.
wi01099485	OSPF adjacency between access switches does not come up over L2 VSN when IGMP snooping is enabled on c-vlan on BEB.

WI reference	Description
wi01099642	The MAC address record for the IS-IS SYS ID should have the RTR Bit set when IPSC is enabled. The bit was turned off even though IPSC is enabled causing all the IPSC pkts to route through the CP.
wi01100357	Queries to unsupported IS-IS MIBs could result in a device crash. The following MIBs are unsupported and, if polled directly or indirectly, can result in a crash: <ul style="list-style-type: none"> isisSummAddrTable (OID-1.3.6.1.3.37.1.1.4) isisISAdjAreaAddrTable (OID-1.3.6.1.3.37.1.6.2) isisISAdjIPAddrTable (OID-1.3.6.1.3.37.1.6.3) isisISAdjProtSuppTable (OID-1.3.6.1.3.37.1.6.4)
wi01100999	Locking the WAN/Cloud uplink port causes the port to NGW to die and tunnels to stop pinging. The port has <code>diag mirror-by-port</code> set on the NGW port for the problem to occur. IPv6 tunnel-to-tunnel stops working when <code>port access-diffserv</code> is set to true.
wi01101007	Disabling and then enabling a line card causes tunnels to stop. Unused cards also cause tunnels to stop if they are reset.
wi01101725	The browse button for the IP redistribution screen in VRF context does not show all VRFs.
wi01101843	Tunneled traffic takes incorrect traffic flow after initial failover of WAN. On the second failover, all tunnel traffic through the failed side is down. IPv6 BGP static route failover gets stuck. On the second failover, it does not work.
wi01102233	When using sub-second timers for ICMP echo-requests to an ERS 8800/8600, ICMP Echo-replies are momentarily delayed or dropped intermittently when removing or inserting a PCMCIA card on the 8692 CPU/SF.  Important: The above issue is fixed. However, Avaya recommends that you download large files to or from a PCMCIA card during off-hours or a maintenance window to minimize impact on the network.
wi01102946	Duplicated L2 Mcast traffic can result after reboot of core switches one by one. This is a result of a timing race condition.
wi01103339	SSH logout traps are incorrectly sent out for a failed SSH login with the wrong user credentials when one or more successful SSH login sessions exist. SSH logout traps are not sent out when multiple SSH sessions are logged out.
wi01103344	The CLI command <code>config sys set info</code> should display the "sendLoginSuccessTrap" enable/disable status.
wi01104447	Variable bindings sequence of <code>rcn(a)NewBfdStatusChange</code> trap are inconsistent with their MIB definitions.

WI reference	Description
wi01106611	The legacy SNMP IpAddress Type, which should be used to hold IPv4 address, is used to hold IPv6 address.
wi01107370	<p>An SPB BEB node can crash with the following signatures: <code>ip_rtdeleteVrf: orec is NULL! RT:%ld.%ld.%ld.%ld</code> or <code>plsbDeleteBmac: routeTree empty, non-zero count</code>. This can happen in the following scenario:</p> <ul style="list-style-type: none"> • Alternative route feature is disabled in the GRT or VRF under consideration. • There is a network/route that is learned via some routing protocol other than ISIS such as RIP, OSPF, BGP, or static. • The same route is learned via SPB/ISIS. SPB adds that route to the RTM and, if ISIS_SPBM route is more preferred, will replace the existing route. • To replace the route, the RTM first deletes that existing "Best Route." • This triggers a notification to other protocols to submit their "Best" Routes." • This trigger to ISIS/SPB, will cause ISIS/SPB to call RTM to add the same route again in a recursive manner, which can cause corruption of the RTM data structures. • The crash can occur the next time the BEB tries to remove that route.
wi01120076	IPv6 in IPv4 (6-in-4) tunnel packets are forwarded to a subnet where they are not supposed to be seen.
wi01120004	The GET-BULK operation to <code>RcVrrpTmpRouterStatisticsTable</code> fails and the SNMP Agent responds with a TIMEOUT error message when you set the max-repetition to greater than 48 for two entries or 27 for one entry.
wi01123818	Core file is generated after logging out of an SSH or Telnet session.
wi01124332	Display freezes on command: <code>ipv6 interface icmpstats gigabitEthernet</code> .

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.


 **Caution:**

Proper handling of compact flash cards and modules can eliminate many potential issues. Please refer to the following sections to avoid unnecessary problems:

- [Proper care of external compact flash and PCMCIA cards](#) on page 30
- [Proper handling of SF CPU and I O modules](#) on page 31

Release 7.2.10 known issues

Table 13: Release 7.2.10 known issues

WI reference	Description
wi01121249	<p>The switch does not accept “destination IP ACL” configurations that use a multicast IP address as the filter criteria. For example, <code>filter acl 1 ace 1 ip dst-ip eq 239.255.255.250</code> fails and causes an <code>Error: Invalid IP address message</code>.</p> <p>Workaround: Use a custom ACE to define a pattern match for the same multicast IP address. For example, to match the <code>239.255.255.250</code> value used above, use <code>filter acl 1 ace 1 advanced custom-filter1 xyz eq effffffa</code>.</p> <p> Important:</p> <p>During an upgrade, if the system configuration file includes a <code>dst-ip</code> ACL with a multicast IP address, that ACL does not carry over and the configuration trying to load fails.</p>
wi01052269	<p>The <code>show fulltech</code> command stops displaying information when the <code>show ip bgp route-vpnv4</code> command starts to display. This happens when <code>cli more false</code> is enabled in CLI or <code>terminal more disable</code> is set in ACLI.</p>
wi01068569	<p>In IS-IS, if you disable redistribution and then apply a policy, you receive a warning that you need to apply the policy even though you already did.</p>

Known issues and limitations

WI reference	Description
	When you enable redistribution and then apply the policy, you do not receive the warning because you already applied the policy. This is working as expected.
wi01087195	IPv6 address 0::0 can be the static-route next-hop, but cannot be the BFD next-hop.
wi01089689	When using sub-second timers for ICMP echo-requests to an ERS 8800/8600, ICMP Echo-replies are momentarily delayed or dropped when inserting an I/O module, or enabling a previously disabled I/O module, or when running commands like <code>config ipv6 forwarding enable (or disable)</code> .
wi01089695	Inserting or removing a link that is part of a LAG using min-link between ERS 8800/8600 and Passport 15k may cause CPU utilization on ERS 8800 to go to 100% and cause the system to become unresponsive to CLI and SNMP interaction. Workaround: Disable and then re-enable the problem port at the remote or local end of the link to resolve the problem. Because this is an intermittent problem, you may need to repeat this procedure until the problem is resolved.
wi01092914	Autonegotiation between some modules and the Passport 15000 (PP15K) may fail intermittently when the PP15K is rebooted or reset. The modules are: 8634XGRS, 8834XG, 8648GBRS, and 8848GB. Workaround: Turn off autonegotiation and set the appropriate line speed for ports on both ends of the link.
wi01092918	Lack of SSH session identifier makes it impossible to correlate traps for multiple concurrent SSH sessions from same IP address/Host.
wi01092921	SSH traps do not contain the user's login name, which would enable the security administrator to know who tried to login to the ERS but failed.
wi01101099	The commands under the <code>show ipv6 ospf tree</code> do not perform a consistency check before displaying their output.
wi01105152	Traps are not sent out for successful/failed login methods such as Telnet, console, modem, and rlogin. They are sent out only for a successful/failed SSH login and a successful/failed HTTPS EDM login.
wi01106859	SSH unauthorized access and authentication failure traps are not sent when the passphrase is wrong or missing with DSA/RSA authentication.
wi01106863	DSA Authentication, with the password authentication parameter (<code>pass-auth</code>) set to false, is falling back to Password Authentication Mode instead of exiting when there is no passphrase or after three attempts with wrong passphrases.
wi01106885	RSA/DSA Authentication with SSH is only prompting once for a passphrase When you do not enter any passphrase.

WI reference	Description
wi01108722	Disabling and then re-enabling a line card sometimes causes an ECMP error such as COP-SW-IP ERROR Slot 1: ercdProcIpRecMsg: Failed to Add ECMP IP Record. This error is intermittent and the line card recovers on its own.
wi01112779	The rcVrrpExtTrapStateTransition and rcVrrpTmpTrapNewMaster Traps definitions do not have associated InetAddressType object for the InetAddress objects.
wi01111394	Standard bgpPeerTable is not supported in 7.2.10, the status of its MIB definitions should be changed to 'obsolete' or deprecated.

Previously reported known issues


The following sections list known issues in Ethernet Routing Switch 8800/8600 reported in software releases prior to Release 7.2. These may be resolved in a future release.

Platform known issues

Table 14: Platform known issues

WI reference	Description
wi00506367	Line RDI is not generated properly as a result of LOS on the 8683XZR module in WAN mode.
wi00507721	The md5 command cannot be run on the standby CPU.
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.
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.

WI reference	Description
wi00518565	When upgrading FPGA firmware on R or RS modules, the following message can appear: 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 This message has no negative effect on the FPGA upgrade. There are no specific FPGA upgrades required with release 7.0.
wi00518661	If you enable DDM monitoring on a switch with non-DDM GBICs installed, the switch generates a message (HAL INFO GBIC) every 5 seconds to the console and to the log file for each non-DDM GBIC installed.
wi00518696	For the system power supply calculation, a low inaccurate value (410 W) is associated with any power supply that displays as unrecognized . 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 unrecognized .
wi00824070	After rebooting the chassis several times in quick succession, the chassis may become unresponsive.
wi00824072	After a switch boot, if you attempt to execute the format /pcmcia command immediately, the console may become unresponsive. Workaround: Wait for the system to complete the boot process and write logs before trying to format the flash.
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.
wi00908084	Random Blank "SW WARNING" messages might be seen. These are not service impacting and can be ignored.
wi00923642	When using scripts to generate repeated access to the file system, the Telnet or console session may become unresponsive as the file system attempts to make multiple updates.
wi00968806	ERS 10 Gbps modules (8612XLRS and 8812XL) do not always apply ACL filters actions such as deny and count on a 10 Gbps port.
wi00972011	When slots are reset during HA, the backup CP may attempt to access records that are not present on the backup CP. This may cause an error message that is non service impacting and can safely be ignored because the errors are coming from the disabled slot.
wi00972964	When the number of files stored on the PCMCIA card exceeds 100, the I/O modules may continuously reset after rebooting the chassis. If there are more than 100 files on the PCMCIA card, delete any unnecessary files and reboot the chassis.
wi00995926	In a rare event, the use of FTP to the flash may corrupt files. To recover, format the flash.

WI reference	Description
wi00999951	<p>COP-SW-ERROR and COP-SW-WARNING messages display when you enable copper SFPs that do not have a cable plugged in. Workaround: For copper SFPs inserted into the fiber ports of the 8634XGRS, 8834XG, 8648GBRS, or 8848GB modules, Avaya recommends that you either attach a copper cable to a viable peer so that a link can be established or administratively disable the SFPs.</p> <p> Note: Failure to adhere to this recommendation causes the modules to consume an excessive amount of CPU cycles, which can adversely effect performance and operation of the module.</p>
wi01011391	<p>In a dual CPU switch, the filtered multicast streams in the IGMP sender table are not synchronized between the Primary and Secondary CPUs. Only the Primary CPU maintains the current list of entries.</p>
wi01016194	<p>Making configuration changes too soon after an HA failover may result in loss. Workaround: After the HA event is complete, monitor the CP usage for a short time to ensure the utilization is back to normal operational levels, then proceed with configuration.</p>
wi01020627	<p>Redistributing direct routes between VRFs does not work unless you enable a routing protocol (BGP, OSPF, or RIP). This adds the routes to the routing table manager (RTM) of the destination VRF and re-advertises the routes to the destination protocol.</p>
wi01024877	<p>After HA, the error message <code>ercdReplaceIpmcRecord: IPMC rcdRadixLookup failed</code> sometimes displays and may cause traffic loss.</p>
wi01034243	<p>The ERS 8800/8600 marks the status of a static-rp with passive interfaces as <code>valid</code>. This is incorrect. A static-rp entry that has a route in the routing table, and the outgoing interfaces for that route is passive, should be marked <code>invalid</code>. This is a display issue with no functional impact.</p>
wi01035278	<p>If a record does not exist, you may see the following error message <code>ercdArpClearActivityBit: ARP rcdRadixLookup failed</code>. There is no functional impact so you can ignore this message.</p>
wi01036254	<p>On reboot, the switch displays <code>IPMC ERROR ipmSysDeleteSession</code> error messages when the switch tries to delete sessions that were already deleted. There is no functional impact so you can ignore this message.</p>
wi01037317	<p>After a slot reset, you may see a <code>COP-SW-MGID ERROR Slot 2: ercdDeleteEgressPepRecord: Lookup Failed!</code> message. This happens when the CPU tries to clean up IPMC records that were already deleted. There is no functional impact so you can ignore this message.</p>

WI reference	Description
wi01037823	After a slot reset, you may see <code>IPMC WARNING updateLogicalPhyPortMapIpmcPep</code> message. This happens when the CPU tries to clean up the record for a remote VLAN that does not exist.
wi01038092	After adding an ACL filter, you may see a <code>COP-SW ERROR SyncTransportLayer: error</code> message. There is no functional impact so you can ignore this message.

Switch management known issues


Table 15: Switch management known issues

WI reference	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 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 16: KHI known issues

WI reference	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>
wi00963449	Upgrading other modules causes KHI errors on the 8812XL SFP+ I/O module and potential traffic loss.

WI reference	Description
	Workaround: Take the 8812XL SFP+ I/O module out of service before you upgrade other modules.
wi00700896	<p>KHI error messages are sometimes displayed after rebooting a switch in ACLI mode. These messages appear intermittently and do not cause any traffic loss.</p> <p> Note: You may also see log error messages after making configuration changes. These INFO messages appear only for a short time while the configuration is getting updated.</p>

SPBM known issues

Table 17: SPBM known issues

WI reference	Description
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) (NN46205–525)</i> .
wi00967803	Ethernet Routing Switch 8800/8600 does not support NLB-multicast and NLB-multicast with IGMP for SPBM.
wi01008560	In IP Multicast over SPBM, if you delete any ssm-map in a static range group, the switch deletes the entire static range group. For example, create an ssm-map for 232.122.122.122 to 232.122.122.128. Then configure this same range in a static group. If you delete any ssm-map between 232.122.122.122 to 232.122.122.128, the entire static range group will be deleted.
wi01019782	Multicast over SPBM does not support adding static groups to an L2 VSN with IGMPv3.
wi01036614	Re-creating a Data I-SID can take up to several minutes in the secondary SMLT peer for a multicast stream. This happens only after a double failure of the SMLT BEB where the BEB reboots twice within 10-15 mins.
wi01037018	PIM is not supported on SPBM VLANs even though the switch allows you to enable PIM.
wi01038097	The L1 DR priority is only for broadcast interfaces, which is not supported in this release. The current implementation of IS-IS supports point-to-point interfaces <i>only</i> . Therefore, the default value displays as 0, not 64.

Layer 2 known issues

Table 18: Layer 2 known issues

WI reference	Description
wi00506797	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.

MLT/SMLT known issues

Table 19: MLT/SMLT known issues

WI reference	Description
wi00852923	Enabling VLAN Manager trace level 4 is not recommended. It is recommended to trace directly to the Ethernet port.
wi00934656	LACP is not supported on SMLT NNIs because it can cause traffic to be silently dropped.
wi00958282	Before you decommission an SMLT switch, disable the ports first to prevent loops in the network.
wi00973586	IPv6 issues with SMLT: <ul style="list-style-type: none"> • L2 VSNs in an SMLT environment do not support IPV6 routing configured on a VLAN/interface. • IP shortcuts in an SMLT environment do not support IPV6 routing.
wi01030203	Ports in the same MLT can be applied to one non-default QoS egress-queue template <i>only</i> . This means that you cannot put ports from the same MLT but with different queues into different templates. The only exception is for the default templates.

Unicast routing known issues

Table 20: Unicast routing known issues

WI reference	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.
wi00971374	Ping an OSPFv3 next hop that is a link-local address with a different scope-id and the following error displays <code>RCIP6 ERROR rcip6RpcOutChangeResEntrySubCid: REPLACE neighbor to HW FAILED</code> .
wi01037106	Ping to IPv6 address with <code>source</code> option fails.

CLI and ACLI known issues

Table 21: CLI and ACLI known issues

WI reference	Description
wi00506209	The <code>copy</code> command does not work properly with FTP debug turned on.
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> .
wi01014724	In the ACLI, the timeout for a Telnet session is applied to all open sessions even if there is activity on some sessions. Workaround: Use the CLI, which works as expected.
wi01036682	In the ACLI, the <code>show ip pim interface</code> and <code>show ip pim neighbour</code> tables do not display the total entry count (count summary).
wi01038428	In the ACLI, the <code>clear ip route <gig></code> command does not work. Workaround: Use the CLI command, which works as expected.

Enterprise Device Manager known issues

Table 22: Enterprise Device Manager known issues

WI reference	Description
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 1016 1339 1094">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>
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 data-bbox="509 1598 1227 1766" 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. 4. Select IP > VRF > Insert to create a VRF.

WI reference	Description
	<p>5. Select IP > IPVPN > VPN > Insert to create an IPVPN for the VRF just created.</p> <p>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.</p>
wi00882843	SNMP table traversals with bulk requests is slower than traversing the same table with Get-Next requests.
wi00970250	EDM copy feature does not work if pasting to another application or another EDM instance.
wi01034797	In EDM, if you do not log out before you close the Browser, the user tables do not get cleaned up. Workaround: Log out of EDM before closing the Browser.
wi01038122	In EDM, the DrHelloTimer units are displayed in milliseconds when they should be in seconds. Workaround: Divide the EDM displayed value for the DrHelloTimer by 1000.

IGMP known issues


Table 23: Release 7.2.10 known issues

WI reference	Description
wi01035454	IGMP interface must be active (with multicast protocol configured) before you can add or delete IGMP static groups.

Unified Access known issues

Table 24: Release 7.2.10 known issues

WI reference	Description
wi00958509	In Unified Access networks, mobility tunnels flap when LACP-enabled SMLT ports are disabled and then enabled. Workaround: The above happens only when group STP is disabled. Avaya recommends enabling STP (group-stp-enabled) when LACP is configured on SMLT ports.
wi00970655	In Unified Access networks, EDM does not provide any indication when the management of a WSP is redirected from one WCP controller to another WCP controller.

WI reference	Description
wi00976807	<p>In Unified Access networks, if a WSP is part of an SMLT Switch Cluster, Avaya recommends configuring the same VLAN server for the WSP and its SMLT peer.</p>
wi00980612	<p>In Unified Access networks, the AP and WSP send WLAN control traffic that is exchanged between them with the high priority bit set. This is to ensure that the control traffic is not dropped even if there is line rate traffic flowing on the port between the AP and the WSP. By default, the ports on the ERS 8800/8600 where the WSP resides are untrusted so they discard the priority bit. This may result in loss of control traffic and bring down the tunnel.</p> <p> Important:</p> <p>Avaya strongly recommends that you configure any port that carries control traffic as trusted so that the priority bits are honored.</p>
wi00984244	<p>In Unified Access networks, you must set the maximum transmission unit (MTU) size on all PoE switches to 1572. If the MTU size of 1572 is not supported, you must enable jumbo frames on all the PoE switches.</p>
wi00987844	<p>In Unified Access networks, when an I/O module goes offline for any reason such as a reset or administratively disabled, the statistics for all WSP tunnels are reset to zero.</p>
wi00995800	<p>In some Unified Access configurations, multicast traffic is sent to a WSP on which a Remote VLAN (RVlan V2) resides, and the wireless clients are registered as receivers for the multicast stream through RVlan V2, the multicast traffic must traverse this WSP twice to get to the receivers.</p> <ul style="list-style-type: none"> • The first pass is from the Sender through WSP1 to the mobility VLAN server (Mvlan Server) on WSP2. • The second pass is from Mvlan Server back through the MT to WSP1 and then through the AT to the mobility clients. <p>As a result, two IpmcPep records are required to handle each of these multicast streams. In this scenario, ERS 8800/8600 switches enabled with Unified Access support up to 3000 IP multicast streams. If this scenario is also in an SMLT setup, four IpmcPep records are required and the ERS 8800/8600 switches support up to 1500 IP multicast streams.</p>
wi01038763	<p>In Unified Access networks, there is a potential for high priority queue drops during a WSP failover.</p> <p>Workaround: Modify the QoS template to provide additional bandwidth to high-priority WLAN traffic. For information on how to make the modification, see the QoS templates section in <i>Configuration — Unified Access</i> (NN46205–526).</p>