NOKIA

Service Router | Release 15.1.R3

SR OS Software Release Notes

3HE 13648 0003 TQZZA 01

Issue: 01

March 2018

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

Table of Contents

List of tables

1 Release Notice

1.1 About this Document

This document provides an overview of the Service Router Operating System (SR OS) in Release 15.1.R3 for the 7450 Ethernet Service Switch (ESS), 7750 Service Router (SR), and 7950 eXtensible Routing System (XRS) platforms.

1.2 Release 15.1.R3 Documentation Set

The SR OS Release 15.1.R3 documentation set consists of Release Notes and the 7450 ESS, 7750 SR, and 7950 XRS user guides. The components of the documentation set are listed in Table 1. New guides introduced in Release 15.1 are highlighted in **bold**.

➡ **Note:** Starting with Release 14.0.R1, the product documentation for the 7450 ESS, 7750 SR, and 7950 XRS platforms has been combined into one documentation set.

Table 1 **Release 15.1.R3 Documentation Set**

Document title	Platform	Part number
SR OS 15.1.R3 Software Release Notes	7450 ESS 7750 SR 7950 XRS	3HE 13648 0003 TQZZA
SR OS 15.1 AA Protocols and Applications	7450 ESS 7750 SR	3HE 13723 0000 TQZZA
Acronyms Reference Guide	7450 ESS 7750 SR 7950 XRS	3HE 13627 AAAA TQZZA
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 15.0.R5 - Part I	7450 ESS 7750 SR 7950 XRS	3HE 13717 AAAA TQZZA

Table 1 **Release 15.1.R3 Documentation Set (Continued)**

Document title	Platform	Part number
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 15.0.R5 - Part II	7450 ESS 7750 SR 7950 XRS	3HE 13718 AAAA TQZZA
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 15.0.R5 - Part III	7450 ESS 7750 SR 7950 XRS	3HE 13719 AAAA TQZZA
Basic System Configuration Guide	7450 ESS 7750 SR 7950 XRS	3HE 13628 AAAA TQZZA
Documentation Suite Overview	7450 ESS 7750 SR 7950 XRS	3HE 13629 AAAA TQZZA
Interface Configuration Guide	7450 ESS 7750 SR 7950 XRS	3HE 13630 AAAA TQZZA
Gx AVPs Reference Guide	7750 SR	3HE 13631 AAAA TQZZA
Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN	7450 ESS 7750 SR 7950 XRS	3HE 13632 AAAA TQZZA
Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services	7450 ESS 7750 SR 7950 XRS	3HE 13633 AAAA TQZZA
Log Events Guide	7450 ESS 7750 SR 7950 XRS	3HE 13634 AAAA TQZZA
MPLS Guide	7450 ESS 7750 SR 7950 XRS	3HE 13635 AAAA TQZZA
Multicast Routing Protocols Guide	7450 ESS 7750 SR 7950 XRS	3HE 13636 AAAA TQZZA
Multiservice Integrated Service Adapter Guide	7450 ESS 7750 SR	3HE 13637 AAAA TQZZA

Table 1 **Release 15.1.R3 Documentation Set (Continued)**

Document title	Platform	Part number
OAM and Diagnostics Guide	7450 ESS 7750 SR 7950 XRS	3HE 13638 AAAA TQZZA
Quality of Service Guide	7450 ESS 7750 SR 7950 XRS	3HE 13639 AAAA TQZZA
RADIUS Attributes Reference Guide	7750 SR	3HE 13640 AAAA TQZZA
Router Configuration Guide	7450 ESS 7750 SR 7950 XRS	3HE 13641 AAAA TQZZA
Services Overview Guide	7450 ESS 7750 SR 7950 XRS	3HE 13642 AAAA TQZZA
System Management Guide	7450 ESS 7750 SR 7950 XRS	3HE 13643 AAAA TQZZA
Triple Play Service Delivery Architecture Guide	7450 ESS 7750 SR	3HE 13644 AAAA TQZZA
Troubleshooting Guide	7450 ESS 7750 SR	3HE 11475 AAAB TQZZA
Unicast Routing Protocols Guide	7450 ESS 7750 SR 7950 XRS	3HE 13645 AAAA TQZZA
Zipped collection of documents	7450 ESS 7750 SR 7950 XRS	3HE 13646 AAAA TQZZA

Note: The *Versatile Service Module Guide* is no longer delivered as a part of the SR OS customer documentation suite. The information related to this module is now included in the *7450 ESS, 7750 SR, and 7950 XRS Basic System Configuration Guide*.

1.3 Guide Conventions

This guide uses the following terminology:

- SR OS node, SR OS chassis—the 7450 ESS, 7750 SR, and 7950 XRS platforms
- NFM-P—the IP/MPLS network and service management functions of the Nokia 5620 Service Aware Manager (SAM) in Release 15.0 and later. The 5620 SAM is now incorporated as part of the Network Services Platform (NSP) software package as the Network Functions Manager for Packet (NFM-P) module.
- ISA—any of the following hardware assemblies, unless otherwise stated:
  - MS-ISA/MS-ISA-E cards (for example, the 7750 SR/7450 ESS Multiservice ISA)
  - MS-ISM/MS-ISM-E line cards
  - any IMMs containing MS-ISA2/MS-ISA2-E cards (for example, 7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3BQ)
  - MS-ISA2 and MS-ISA2-E in IOM4-e and 7750 SR-e

2 Release 15.1.R3 Supported Hardware

The following tables summarize the hardware supported in SR OS Release 15.1.R3. New hardware supported in SR OS Release 15.1 is printed in **bold**.

2.1 Supported Chassis Configurations

Table 2 **Supported 7950 XRS Chassis Configurations**

Nokia Model #	Description
7950 XRS-16c	A single 33RU chassis that holds up to 8 XCMs and 16 C-XMAs
7950 XRS-20	A single 44RU chassis (when equipped with optional hood rear air deflector) that holds up to 10 XCMs and 20 XMAs or C-XMAs
7950 XRS-40	Comprised of two XRS-20 chassis or two XRS-20e chassis that can hold up to 40 XMAs
7950 XRS-20e Universal Chassis	A single 44RU chassis that holds up to 10 XCMs and 20 XMAs or C-XMAs. The XRS-20e Universal chassis supports any of Low Voltage DC (LVDC), AC, and HVDC power options.
7950 XRS-20e AC/HVDC	A single 44RU chassis that holds up to 10 XCMs and 20 XMAs or C-XMAs. The XRS-20e AC/HVDC chassis supports any of AC or HVDC power options.

Table 3 **Supported 7750 SR and 7450 ESS Chassis**

Nokia Model #	Description
7750 SR-7	7750 SR-7 chassis (DC; AC requires external AC Rectifier shelf)
7750 SR-7-B	7750 SR-7-B chassis (DC; AC requires external AC Rectifier shelf)
7750 SR-12	7750 SR-12 chassis (DC; AC requires external AC Rectifier shelf)
7750 SR-12-B	7750 SR-12-B chassis (DC; AC requires external AC Rectifier shelf)
7750 SR-12e	7750 SR-12e integrated chassis
7750 SR-a4	7750 SR-a4 chassis (AC and DC)
7750 SR-a8	7750 SR-a8 chassis (AC and DC)
7750 SR-c4	7750 SR-c4 chassis (AC and DC)

Table 3 **Supported 7750 SR and 7450 ESS Chassis (Continued)**

Nokia Model #	Description
7750 SR-c12	7750 SR-c12 chassis (AC and DC)
7750 SR-1	**7750 SR-1 chassis (AC and DC)**
7750 SR-1e	7750 SR-1e chassis (AC and DC)
7750 SR-2e	7750 SR-2e chassis (AC and DC)
7750 SR-3e	7750 SR-3e chassis (AC and DC)
7450 ESS-7	7450 ESS-7 chassis (AC and DC)
7450 ESS-12	7450 ESS-12 chassis (AC and DC)

2.2 Supported Cards (SFM, CPM, XCM, CCM, CFM, MCM, IOM, IMM, ISM)

The following tables summarize the Switch Fabric/Control Processor Modules (SF/CPMs, CPMs, or SFMs), XMA Control Modules (XCMs), Connection and Control Modules (CCMs), Control and Forwarding Modules (CFMs), MDA Carrier Modules (MCMs), Chassis Control Modules (CCMs), Input/Output Modules (IOMs), Integrated Media Modules (IMMs), and Integrated Services Modules (ISMs) supported in SR OS.

Table 4 **SFM, CPM, CCM, and XCM Cards Supported in 7950 XRS**

Nokia Part #	Description	CLI String (Card)
3HE06936AA	7950 XRS-20 XMA Control Module (XCM-X20)	xcm-x20
3HE07115AA	7950 XRS-20 Switch Fabric Module (SFM-X20)	sfm-x20
3HE07116AA	7950 XRS-20 Control Processor Module (CPM-X20)	cpm-x20
3HE07116AB	7950 XRS-20 Control Processor Module (CPM-X20) 16GB	cpm-x20
3HE07117AA	7950 XRS-20 Connection and Control Module (CCM-X20)	ccm-x20
3HE08021AA	7950 XRS-20 Switch Fabric Module B (SFM-X20-B)	sfm-x20-b
3HE08120AA	7950 XRS-16c Switch Fabric Module (SFM-X16)	sfm-x16-b

Table 4 **SFM, CPM, CCM, and XCM Cards Supported in 7950 XRS**

Nokia Part #	Description	CLI String (Card)
3HE08121AA	7950 XRS-16c Control Processor Module (CPM-X16)	cpm-x16
3HE08125AA	7950 XRS-16c XMA Control Module (XCM-X16)	xcm-x16
3HE08505AA	7950 XRS-20 Standalone Switch Fabric Module B (SFM-X20S-B)	sfm-x20s-b
3HE09280AA	7950 XRS-16c XCM with XMA support	xcm-x16
3HE11087AA	XCM - 7950 XRS-20e XCM	xcm-x20

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE03607AA	7750 SR-c12 CFM-XP	✓						cfm-xp	--
3HE03608AA	7750 SR-c4/c12 MCM-XP[1]	✓						mcm-xp	--
3HE03617AA	7750 SR-12 SF/CPM3					✓		sfm3-12	--
3HE03619AA	7750 SR IOM3-XP				✓	✓	✓	iom3-xp	--
3HE03622AA	7750 SR 4-port 10GE XFP IMM				✓	✓		imm4-10gb-xfp	imm2-10gb-xp-xfp imm2-10gb-xp-xfp
3HE03623AA	7750 SR 8-port 10GE XFP IMM				✓	✓		imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE03624AA	7750 SR 48-port GE SFP IMM				✓	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE03625AA	7750 SR 48-port GE copper/TX IMM				✓	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE04164AA	7750 SR-7 SF/CPM3				✓			sfm3-7	--
3HE04580AA	7750 SR-c12 CCM-XP	✓						ccm-xp	--
3HE04741AA	7750 SR 5-port 10GE XFP IMM				✓	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE04743AA	7x50 12-port 10G Ethernet SFP+ IMM – L3HQ				✓	✓		imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05053AA	7x50 1-port 100G Ethernet CFP IMM – L3HQ				✓	✓		imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05553AA	7x50 12-port 10G Ethernet SFP+ IMM – L2HQ				✓	✓		imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05553BA	7x50 12-port 10G Ethernet SFP+ IMM – L3BQ				✓	✓		imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05814AA	7x50 1-port 100G Ethernet CFP IMM – L2HQ				✓	✓		imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05814BA	7x50 1-port 100G Ethernet CFP IMM – L3BQ				✓	✓		imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05895AA	7x50 48-port GE SFP IMM – L2HQ				✓	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05895BA	7x50 48-port GE SFP IMM – L3BQ				✓	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05896AA	7x50 48-port GE copper/ TX IMM – L2HQ				✓	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05896BA	7x50 48-port GE copper/ TX IMM – L3BQ				✓	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05898AA	7x50 5-port 10GE XFP IMM – L2HQ				✓	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05898BA	7x50 5-port 10GE XFP IMM – L3BQ				✓	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05899AA	7x50 8-port 10GE XFP IMM – L2HQ				✓	✓		imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE05899BA	7x50 8-port 10GE XFP IMM – L3BQ				✓	✓		imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE05948AA	7750 SR-12 SF/CPM4					✓		sfm4-12	--
3HE05949AA	7750 SR-7 SF/CPM4				✓			sfm4-7	--
3HE06318AA	7750 Multicore-CPU IOM3-XP-B				✓	✓	✓	iom3-xp-b	--
3HE06320AA	7x50 3-port 40GE QSFP IMM – L3HQ				✓	✓		imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06326AA	7x50 48-port GE Multicore-CPU SFP IMM – L3HQ				✓	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326BA	7x50 48-port GE Multicore-CPU SFP IMM – L3BQ				✓	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326CA	7x50 48-port GE Multicore-CPU SFP IMM – L2HQ				✓	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06428AA	7x50 48-port GE SFP IMM – L3HQ				✓	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06429AA	7x50 48-port GE copper/ TX IMM – L3HQ				✓	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE06430AA	7x50 5-port 10GE XFP IMM – L3HQ				✓	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE06431AA	7x50 8-port 10GE XFP IMM – L3HQ				✓	✓		imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE06721AA	7x50 3-port 40GE QSFP IMM – L2HQ				✓	✓		imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06721BA	7x50 3-port 40GE QSFP IMM – L3BQ				✓	✓		imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE07158AA	7x50 12-port 10GE FP3 SFP+ IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE07158BA	7x50 12-port 10GE FP3 SFP+ IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158CA	7x50 12-port 10GE FP3 SFP+ IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07159AA	7x50 1-port 100GE FP3 CFP IMM – L3HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07159BA	7x50 1-port 100GE FP3 CFP IMM – L3BQ				✓	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07159CA	7x50 1-port 100GE FP3 CFP IMM – L2HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07166AA	7750 SR-12e SF/CPM4-12e						✓	sfm4-12e	--
3HE07167AA	7750 SR-12e Mini-SFM4-12e						✓	m-sfm4-12e	--
3HE07303AA	7x50 2-port 100GE FP3 CFP IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303BA	7x50 2-port 100GE FP3 CFP IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303CA	7x50 2-port 100GE FP3 CFP IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07304AA	7x50 6-port 40GE FP3 QSFP IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304BA	7x50 6-port 40GE FP3 QSFP IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304CA	7x50 6-port 40GE FP3 QSFP IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07305AA	7x50 20-port 10GE FP3 SFP+ IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE07305BA	7x50 20-port 10GE FP3 SFP+ IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp

Table 5 SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE07305CA	7x50 20-port 10GE FP3 SFP+ IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE08019AA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08019BA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3BQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08019CA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L2HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08020AA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020BA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020CA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08173AA	7750 SR-c12 CFM-XP-B	✓						cfm-xp-b	--
3HE08174AA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174BA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174CA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08175AA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE08175BA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175CA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08421AA	7750 SR SF/CPM5-12e						✓	sfm5-12e	--
3HE08422AA	7750 SR Mini-SFM5-12e						✓	m-sfm5-12e	--
3HE08423AA	7750 SR CPM5				✓	✓	✓	cpm5	--
3HE08424AA	7x50 40-port 10GE SFP+ IMM – L3HQ						✓	imm40-10gb-sfp	m40-10g-sfp
3HE08424BA	7x50 40-port 10GE SFP+ IMM – L3BQ						✓	imm40-10gb-sfp	m40-10g-sfp
3HE08424CA	7x50 40-port 10GE SFP+ IMM – L2HQ						✓	imm40-10gb-sfp	m40-10g-sfp
3HE08425AA	7x50 4-port 100GE CXP IMM – L3HQ						✓	imm4-100gb-cxp	m4-100g-cxp
3HE08425BA	7x50 4-port 100GE CXP IMM – L3BQ						✓	imm4-100gb-cxp	m4-100g-cxp
3HE08425CA	7x50 4-port 100GE CXP IMM – L2HQ						✓	imm4-100gb-cxp	m4-100g-cxp
3HE08426AA	7750 SR IOM3-XP-C				✓	✓	✓	iom3-xp-c	--
3HE08428AA	7750 SR SFM5-12					✓		sfm5-12	--
3HE08429AA	7750 SR SFM5-7				✓			sfm5-7	--
3HE09117AA	7x50 Multiservice ISM				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p-isa2-ms
3HE09118AA	7x50 Multiservice ISM-E (no encryption)				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p-isa2-ms-e

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE09192AA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192BA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192CA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193AA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193BA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193CA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09201AA	7750 SR-a CPM		✓					cpm-a	--
3HE09202AA	7750 SR-a IOM – L3HQ		✓					iom-a	--
3HE09202BA	7750 SR-a IOM – L3BQ		✓					iom-a	--
3HE09202CA	7750 SR-a IOM – L2HQ		✓					iom-a	--
3HE09253AA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253BA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253CA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE09254AA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L3HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254BA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L3BQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254CA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L2HQ				✓	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09279AA	7x50 48-port GE Multicore SFP IMM – L3HQ				✓	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279BA	7x50 48-port GE Multicore SFP IMM – L3BQ				✓	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279CA	7x50 48-port GE Multicore SFP IMM – L2HQ				✓	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09436AA	IMM – 7750 SR 1-PT 100GE INT DWDM L3HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436BA	IMM – 7750 SR 1-PT 100GE INT DWDM L3BQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436CA	IMM – 7750 SR 1-PT 100GE INT DWDM L2HQ				✓	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09645AA	7x50 4-Port 100GE CFP4 IMM – L3HQ						✓	imm4-100gb-cfp4	m4-100g-cfp4
3HE09645BA	7x50 4-Port 100GE CFP4 IMM – L3BQ						✓	imm4-100gb-cfp4	m4-100g-cfp4
3HE09645CA	7x50 4-Port 100GE CFP4 IMM – L2HQ						✓	imm4-100gb-cfp4	m4-100g-cfp4

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE09648AA	IOM – 7750 SR IOM4-e L3HQ				✓	✓	✓	iom4-e	--
3HE09648BA	IOM – 7750 SR IOM4-e L3BQ				✓	✓	✓	iom4-e	--
3HE09648CA	IOM – 7750 SR IOM4-e L2HQ				✓	✓	✓	iom4-e	--
3HE10014AA	IMM – 160-port GE cSFP/80-port GE SFP – L3HQ				✓	✓	✓	imm-1pac-fp3	p160-1gb-csfp
3HE10014BA	IMM – 160-port GE cSFP/80-port GE SFP – L3BQ				✓	✓	✓	imm-1pac-fp3	p160-1gb-csfp
3HE10014CA	IMM – 160-port GE cSFP/80-port GE SFP – L2HQ				✓	✓	✓	imm-1pac-fp3	p160-1gb-csfp
3HE10309AA	CCM – 7750 SR-e CCM-e			✓				ccm-e	--
3HE10310AA	CPM – 7750 SR-e CPM-e			✓				cpm-e	--
3HE10311AA	IOM – 7750 SR IOM-e L3HQ			✓				iom-e	--
3HE10311BA	IOM – 7750 SR IOM-e L2HQ			✓				iom-e	--
3HE10311CA	IOM – 7750 SR IOM-e L3BQ			✓				iom-e	--
3HE10717AA	IOM - 7750 SR IOM4-e-B L3HQ				✓	✓	✓	iom4-e-b	--
3HE10717BA	IOM - 7750 SR IOM4-e-B L3BQ				✓	✓	✓	iom4-e-b	--
3HE10717CA	IOM - 7750 SR IOM4-e-B L2HQ				✓	✓	✓	iom4-e-b	--

Table 5 **SFM, CPM, CFM, MCM, CCM, IOM, IMM, and ISM Cards Supported in 7750 SR**

Nokia Part #	Description	SR-c12	SR-a4/a8	SR-1e/2e/3e	SR-7	SR-12	SR-12e	CLI String (Card)	CLI String (MDA)
3HE11351AA	**IOM - 7750 SR IOM4-e-HS L3HQ**				✓	✓	✓	iom4-e-hs	--
3HE11351CA	**IOM - 7750 SR IOM4-e-HS L2HQ**				✓	✓	✓	iom4-e-hs	--

Note:

1. The MCM, not MCM-XP, is supported in the 7750 SR-c4.

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE03618AA	7450 ESS-12 SF/CPM3		✓	sfm3-12	--
3HE03619AA	7750 SR IOM3-XP	✓	✓	iom3-xp	--
3HE03620AA	7450 ESS IOM3-XP	✓	✓	iom3-xp	--
3HE03622AA	7750 SR 4-port 10GE XFP IMM	✓	✓	imm4-10gb-xfp	imm2-10gb-xp-xfp imm2-10gb-xp-xfp
3HE03623AA	7750 SR 8-port 10GE XFP IMM	✓	✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE03624AA	7750 SR 48-port GE SFP IMM	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE03625AA	7750 SR 48-port GE copper/TX IMM	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE04166AA	7450 ESS-7 SF/CPM3	✓		sfm3-7	--
3HE04741AA	7750 SR 5-port 10GE XFP IMM	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE04743AA	7x50 12-port 10G Ethernet SFP+ IMM – L3HQ	✓	✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05053AA	7x50 1-port 100G Ethernet CFP IMM- L3HQ	✓	✓	imm1-100gb-cfp	imm1-100gb-xp-cfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE05553AA	7x50 12-port 10G Ethernet SFP+ IMM – L2HQ	✓	✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05553BA	7x50 12-port 10G Ethernet SFP+ IMM – L3BQ	✓	✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05814AA	7x50 1-port 100G Ethernet CFP IMM – L2HQ	✓	✓	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05814BA	7x50 1-port 100G Ethernet CFP IMM – L3BQ	✓	✓	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05895AA	7x50 48-port GE SFP IMM – L2HQ	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05895BA	7x50 48-port GE SFP IMM – L3BQ	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05896AA	7x50 48-port GE copper/TX IMM – L2HQ	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05896BA	7x50 48-port GE copper/TX IMM – L3BQ	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05898AA	7x50 5-port 10GE XFP IMM – L2HQ	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05898BA	7x50 5-port 10GE XFP IMM – L3BQ	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05899AA	7x50 8-port 10GE XFP IMM – L2HQ	✓	✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE05899BA	7x50 8-port 10GE XFP IMM – L3BQ	✓	✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE05950AA	7450 ESS-12 SF/CPM4		✓	sfm4-12	--
3HE05951AA	7450 ESS-7 SF/CPM4	✓		sfm4-7	--
3HE06318AA	7750 Multicore-CPU IOM3-XP-B	✓	✓	iom3-xp-b	--
3HE06320AA	7x50 3-port 40GE QSFP IMM- L3HQ	✓	✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06324AA	7450 Multicore-CPU IOM3-XP-B	✓	✓	iom3-xp-b	--
3HE06326AA	7x50 48-port GE Multicore-CPU SFP IMM – L3HQ	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE06326BA	7x50 48-port GE Multicore-CPU SFP IMM – L3BQ	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326CA	7x50 48-port GE Multicore-CPU SFP IMM – L2HQ	✓	✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06428AA	7x50 48-port GE SFP IMM – L3HQ	✓	✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06429AA	7x50 48-port GE copper/TX IMM – L3HQ	✓	✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE06430AA	7x50 5-port 10GE XFP IMM – L3HQ	✓	✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE06431AA	7x50 8-port 10GE XFP IMM – L3HQ	✓	✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE06721AA	7x50 3-port 40GE QSFP IMM – L2HQ	✓	✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06721BA	7x50 3-port 40GE QSFP IMM – L3BQ	✓	✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE07158AA	7x50 12-port 10GE FP3 SFP+ IMM – L3HQ	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158BA	7x50 12-port 10GE FP3 SFP+ IMM – L3BQ	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158CA	7x50 12-port 10GE FP3 SFP+ IMM – L2HQ	✓	✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07159AA	7x50 1-port 100GE FP3 CFP IMM – L3HQ	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07159BA	7x50 1-port 100GE FP3 CFP IMM – L3BQ	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07159CA	7x50 1-port 100GE FP3 CFP IMM – L2HQ	✓	✓	imm-1pac-fp3	p1-100g-cfp
3HE07303AA	7x50 2-port 100GE FP3 CFP IMM – L3HQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303BA	7x50 2-port 100GE FP3 CFP IMM – L3BQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE07303CA	7x50 2-port 100GE FP3 CFP IMM – L2HQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07304AA	7x50 6-port 40GE FP3 QSFP IMM – L3HQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304BA	7x50 6-port 40GE FP3 QSFP IMM – L3BQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304CA	7x50 6-port 40GE FP3 QSFP IMM – L2HQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07305AA	7x50 20-port 10GE FP3 SFP+ IMM – L3HQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE07305BA	7x50 20-port 10GE FP3 SFP+ IMM – L3BQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE07305CA	7x50 20-port 10GE FP3 SFP+ IMM – L2HQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE08019AA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3HQ	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08019BA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3BQ	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08019CA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L2HQ	✓	✓	imm-1pac-fp3	p1-100g-tun
3HE08020AA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3HQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020BA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3BQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020CA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L2HQ	✓	✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08174AA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3HQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174BA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3BQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE08174CA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L2HQ	✓	✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08175AA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3HQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175BA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3BQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175CA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L2HQ	✓	✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08426AA	7750 SR IOM3-XP-C	✓	✓	iom3-xp-c	--
3HE08427AA	7450 ESS IOM3-XP-C	✓	✓	iom3-xp-c	--
3HE08430AA	7450 ESS SFM5-12		✓	sfm5-12	--
3HE08431AA	7450 ESS SFM5-7	✓		sfm5-7	--
3HE08432AA	7450 ESS CPM5	✓	✓	cpm5	--
3HE09117AA	7x50 Multiservice ISM	✓	✓	imm-2pac-fp3	p-isa2-ms p-isa2-ms
3HE09118AA	7x50 Multiservice ISM-E (no encryption)	✓	✓	imm-2pac-fp3	p-isa2-ms-e p-isa2-ms-e
3HE09192AA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3HQ	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192BA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3BQ	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192CA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L2HQ	✓	✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193AA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3HQ	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193BA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3BQ	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193CA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L2HQ	✓	✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE09253AA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L3HQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253BA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L3BQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253CA	7x50 MS-ISA2-E + 1-port 100GE CFP IMM – L2HQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09254AA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L3HQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254BA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L3BQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254CA	7x50 MS-ISA2-E + 10-port 10G SFP+ IMM – L2HQ	✓	✓	imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09279AA	7x50 48-port GE Multicore SFP IMM – L3HQ	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279BA	7x50 48-port GE Multicore SFP IMM – L3BQ	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279CA	7x50 48-port GE Multicore SFP IMM – L2HQ	✓	✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09436AA	IMM – 7750 SR 1-PT 100GE INT DWDM L3HQ	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436BA	IMM – 7750 SR 1-PT 100GE INT DWDM L3BQ	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436CA	IMM – 7750 SR 1-PT 100GE INT DWDM L2HQ	✓	✓	imm-1pac-fp3	p1-100g-tun-b
3HE09648AA	IOM – 7750 SR IOM4-e L3HQ	✓	✓	iom4-e	--
3HE09648BA	IOM – 7750 SR IOM4-e L3BQ	✓	✓	iom4-e	--
3HE09648CA	IOM – 7750 SR IOM4-e L2HQ	✓	✓	iom4-e	--
3HE10014AA	IMM – 160-port GE cSFP/80-port GE SFP – L3HQ	✓	✓	imm-1pac-fp3	p160-1gb-csfp
3HE10014BA	IMM – 160-port GE cSFP/80-port GE SFP – L3BQ	✓	✓	imm-1pac-fp3	p160-1gb-csfp

Table 6 **SFM, CPM, IOM, IMM, and ISM Cards Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	ESS-7	ESS-12	CLI String (Card)	CLI String (MDA)
3HE10014CA	IMM – 160-port GE cSFP/80-port GE SFP – L2HQ	✓	✓	imm-1pac-fp3	p160-1gb-csfp
3HE10717AA	IOM - 7750 SR IOM4-e-B L3HQ	✓	✓	iom4-e-b	--
3HE10717BA	IOM - 7750 SR IOM4-e-B L3BQ	✓	✓	iom4-e-b	--
3HE10717CA	IOM - 7750 SR IOM4-e-B L2HQ	✓	✓	iom4-e-b	--

Note:

1. The isa-*type* can be isa-bb, isa2-aa, isa2-bb, or isa2-tunnel.

Table 7 summarizes the IOMs, IMMs, and ISMs supported in SR OS for the 7450 ESS in mixed mode.

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE03619AA	7750 SR IOM3-XP	iom3-xp	--
3HE03622AA	7750 SR 4-port 10GE XFP IMM	imm4-10gb-xfp	imm2-10gb-xp-xfp imm2-10gb-xp-xfp
3HE03623AA	7750 SR 8-port 10GE XFP IMM	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE03624AA	7750 SR 48-port GE SFP IMM	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE03625AA	7750 SR 48-port GE copper/TX IMM	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE04741AA	7750 SR 5-port 10GE XFP IMM	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE04743AA	7750 SR 12-port 10G Ethernet SFP+ IMM	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05053AA	7750 SR 1-port 100G Ethernet CFP IMM	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05553AA	7x50 12-port 10G Ethernet SFP+ IMM – L2HQ	imm12-10gb-sf+	imm12-10gb-xp-sf+

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE05553BA	7x50 12-port 10G Ethernet SFP+ IMM – L3BQ	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05814AA	7x50 1-port 100G Ethernet CFP IMM – L2HQ	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05814BA	7x50 1-port 100G Ethernet CFP IMM – L3BQ	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05895AA	7x50 48-port GE SFP IMM – L2HQ	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05895BA	7x50 48-port GE SFP IMM – L3BQ	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05896AA	7x50 48-port GE copper/TX IMM – L2HQ	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05896BA	7x50 48-port GE copper/TX IMM – L3BQ	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05898AA	7x50 5-port 10GE XFP IMM – L2HQ	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05898BA	7x50 5-port 10GE XFP IMM – L3BQ	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05899AA	7x50 8-port 10GE XFP IMM – L2HQ	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE05899BA	7x50 8-port 10GE XFP IMM – L3BQ	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE06318AA	7750 Multicore-CPU IOM3-XP-B	iom3-xp-b	--
3HE06320AA	7x50 3-port 40GE QSFP IMM- L3HQ	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06326AA	7x50 48-port GE Multicore-CPU SFP IMM – L3HQ	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326BA	7x50 48-port GE Multicore-CPU SFP IMM – L3BQ	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326CA	7x50 48-port GE Multicore-CPU SFP IMM – L2HQ	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06428AA	7x50 48-port GE SFP IMM – L3HQ	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06429AA	7x50 48-port GE copper/TX IMM – L3HQ	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE06430AA	7x50 5-port 10GE XFP IMM – L3HQ	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE06431AA	7x50 8-port 10GE XFP IMM – L3HQ	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE06721AA	7x50 3-port 40GE QSFP IMM – L2HQ	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06721BA	7x50 3-port 40GE QSFP IMM – L3BQ	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE07158AA	7x50 12-port 10GE FP3 SFP+ IMM – L3HQ	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158BA	7x50 12-port 10GE FP3 SFP+ IMM – L3BQ	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158CA	7x50 12-port 10GE FP3 SFP+ IMM – L2HQ	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07159AA	7x50 1-port 100GE FP3 CFP IMM – L3HQ	imm-1pac-fp3	p1-100g-cfp
3HE07159BA	7x50 1-port 100GE FP3 CFP IMM – L3BQ	imm-1pac-fp3	p1-100g-cfp
3HE07159CA	7x50 1-port 100GE FP3 CFP IMM – L2HQ	imm-1pac-fp3	p1-100g-cfp
3HE07303AA	7x50 2-port 100GE FP3 CFP IMM – L3HQ	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303BA	7x50 2-port 100GE FP3 CFP IMM – L3BQ	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303CA	7x50 2-port 100GE FP3 CFP IMM – L2HQ	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07304AA	7x50 6-port 40GE FP3 QSFP IMM – L3HQ	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304BA	7x50 6-port 40GE FP3 QSFP IMM – L3BQ	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304CA	7x50 6-port 40GE FP3 QSFP IMM – L2HQ	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07305AA	7x50 20-port 10GE FP3 SFP+ IMM – L3HQ	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE07305BA	7x50 20-port 10GE FP3 SFP+ IMM – L3BQ	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE07305CA	7x50 20-port 10GE FP3 SFP+ IMM – L2HQ	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE08019AA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3HQ	imm-1pac-fp3	p1-100g-tun
3HE08019BA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L3BQ	imm-1pac-fp3	p1-100g-tun
3HE08019CA	7x50 1-port 100GE DWDM Tunable FP3 IMM – L2HQ	imm-1pac-fp3	p1-100g-tun
3HE08020AA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3HQ	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020BA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L3BQ	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020CA	7x50 1-port 100GE CFP + 10-port 10GE SFP+ FP3 IMM – L2HQ	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08174AA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3HQ	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174BA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L3BQ	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174CA	7x50 10-port 10GE SFP+ + 20-port GE SFP FP3 IMM – L2HQ	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08175AA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3HQ	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175BA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L3BQ	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175CA	7x50 3-port 40GE QSFP + 20-port GE SFP FP3 IMM – L2HQ	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08426AA	7750 SR IOM3-XP-C	iom3-xp-c	--
3HE09117AA	7x50 Multiservice ISM [1]	imm-2pac-fp3	p-isa2-ms p-isa2-ms

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE09192AA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3HQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192BA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L3BQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192CA	7x50 MS-ISA2 + 1-port 100GE CFP IMM – L2HQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193AA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3HQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193BA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L3BQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193CA	7x50 MS-ISA2 + 10-port 10GE SFP+ IMM – L2HQ [1]	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09279AA	7x50 48-port GE Multicore SFP IMM – L3HQ	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279BA	7x50 48-port GE Multicore SFP IMM – L3BQ	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279CA	7x50 48-port GE Multicore SFP IMM – L2HQ	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09436AA	IMM – 7750 SR 1-PT 100GE INT DWDM L3HQ	imm-1pac-fp3	p1-100g-tun-b
3HE09436BA	IMM – 7750 SR 1-PT 100GE INT DWDM L3BQ	imm-1pac-fp3	p1-100g-tun-b
3HE09436CA	IMM – 7750 SR 1-PT 100GE INT DWDM L2HQ	imm-1pac-fp3	p1-100g-tun-b
3HE09648AA	IOM – 7750 SR IOM4-e L3HQ	iom4-e	--
3HE09648BA	IOM – 7750 SR IOM4-e L3BQ	iom4-e	--
3HE09648CA	IOM – 7750 SR IOM4-e L2HQ	iom4-e	--
3HE10014AA	IMM – 160-port GE cSFP/80-port GE SFP – L3HQ	imm-1pac-fp3	p160-1gb-csfp
3HE10014BA	IMM – 160-port GE cSFP/80-port GE SFP – L3BQ	imm-1pac-fp3	p160-1gb-csfp

Table 7 **IOM, IMM, and ISM Cards Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (Card)	CLI String (MDA)
3HE10014CA	IMM – 160-port GE cSFP/80-port GE SFP – L2HQ	imm-1pac-fp3	p160-1gb-csfp
3HE10717AA	IOM - 7750 SR IOM4-e-B L3HQ	iom4-e-b	--
3HE10717BA	IOM - 7750 SR IOM4-e-B L3BQ	iom4-e-b	--
3HE10717CA	IOM - 7750 SR IOM4-e-B L2HQ	iom4-e-b	--

Note:

1. MS-ISM and MS-ISA2 applications using MS-ISA2s are not supported in mixed mode with the exception of Application Assurance, IPsec, and NAT. IPsec is not supported with MS-ISM-E and MS-ISA2-E.

2.3 Supported Adapters (XMA, MDA, ISA, CMA, VSM)

The following tables summarize the XRS Media Adapters (XMAs), Media Dependent Adapters (MDAs), Integrated Service Adapters (ISAs), Compact Media Adapters (CMAs), and Versatile Services Modules (VSMs) supported in SR OS.

Table 8 **XMAs and C-XMAs Supported in 7950 XRS**

Nokia Part #	Description	CLI String (MDA)
3HE06937AA	C-XMA – 7950 XRS 20-port 10GE SFP+ – IP Core	cx20-10g-sfp
3HE06938AA	C-XMA – 7950 XRS 2-port 100GE CFP – IP Core	cx2-100g-cfp
3HE06937BA	C-XMA – 7950 XRS 20-port 10GE SFP+ – LSR	cx20-10g-sfp
3HE06938BA	C-XMA – 7950 XRS 2-port 100GE CFP – LSR	cx2-100g-cfp
3HE07297AA	XMA – 7950 XRS 40-port 10GE SFP+ – IP Core	x40-10g-sfp
3HE07297BA	XMA – 7950 XRS 40-port 10GE SFP+ – LSR	x40-10g-sfp
3HE07299AA	XMA – 7950 XRS 4-port 100GE CXP – IP Core	x4-100g-cxp
3HE07299BA	XMA – 7950 XRS 4-port 100GE CXP – LSR	x4-100g-cxp
3HE08214AA	C-XMA – 7950 XRS 6-port 40GE QSFP+ – IP Core	cx6-40g-qsfp

Table 8 **XMAs and C-XMAs Supported in 7950 XRS (Continued)**

Nokia Part #	Description	CLI String (MDA)
3HE08214BA	C-XMA – 7950 XRS 6-port 40GE QSFP+ – LSR	cx6-40g-qsfp
3HE08631AA	C-XMA – 7950 XRS 72-port GE CSFP – IP Core	cx72-1g-csfp
3HE08631BA	C-XMA – 7950 XRS 72-port GE CSFP – LSR	cx72-1g-csfp
3HE08632AA	XMA – 7950 XRS 4-port 100G CFP2 – IP Core	x4-100g-cfp2
3HE08632BA	XMA – 7950 XRS 4-port 100G CFP2 – LSR	x4-100g-cfp2
3HE10677AA	XMA – 7950 XRS 2-PT 100GE INT DWDM – IP Core	x2-100g-tun
3HE10677AB	XMA – 7950 XRS 2-PT 100GE INT DWDM – LSR	x2-100g-tun

Table 9 **MDAs, CMAs, and ISAs Supported in 7750 SR**

Nokia Part #	Description	SR-c4/c12	iom3-xp/-b/-c	iom4-e, iom4-e-b, iom4-e-hs, and SR-1e/2e/3e (iom-e)	SR-a4/a8 (iom-a)	SR-1	CLI String (MDA)
3HE00021AA	60-port 10/100TX MDA – mini-RJ21	✓	✓				m60-10/100eth-tx
3HE00023AA	20-port 100FX MDA – SFP	✓	✓				m20-100eth-sfp
3HE00030AA	1-port 10GBASE-LW/LR MDA with optics – Simplex SC		✓				m1-10gb
3HE00032AA	8-port OC-3c/STM-1c MDA – SFP	✓	✓				m8-oc3-sfp
3HE00033AA	16-port OC-3c/STM-1c MDA – SFP		✓				m16-oc3-sfp
3HE00037AA	8-port OC-12c/STM-4c MDA – SFP		✓				m8-oc12/3-sfp
3HE00038AA	16-port OC-12c/STM-4c MDA – SFP		✓				m16-oc12/3-sfp
3HE00043AA	2-port OC-48c/STM-16c MDA – SFP	✓	✓				m2-oc48-sfp
3HE00044AA	4-port OC-48c/STM-16c MDA – SFP		✓				m4-oc48-sfp

Table 9 **MDAs, CMAs, and ISAs Supported in 7750 SR (Continued)**

Nokia Part #	Description	SR-c4/c12	iom3-xp/-b/-c	iom4-e, iom4-e-b, iom4-e-hs, and SR-1e/2e/3e (iom-e)	SR-a4/a8 (iom-a)	SR-1	CLI String (MDA)
3HE00048AA	1-port OC-192c/STM-64c MDA with SR-1/I-64.1 optic – Simplex SC		✓				m1-oc192
3HE00049AA	1-port OC-192c/STM-64c MDA with IR-2/S-64.2 optic – Simplex SC		✓				m1-oc192
3HE00071AA	4-port ATM OC-12c/STM-4c MDA – SFP	✓	✓				m4-atmoc12/3-sfp
3HE00074AA	16-port ATM OC-3c/STM-1c MDA – SFP		✓				m16-atmoc3-sfp
3HE00709AA	1-port OC-192c/STM-64c MDA with LR-2/L-64.2 optic – Simplex SC		✓				m1-oc192
3HE01020AA	8-port Channelized DS1/E1 CMA – RJ48c	✓					c8-chds1
3HE01021AA	4-port DS3/E3 CMA – 1.0/2.3	✓					c4-ds3
3HE01022AA	8-port 10/100TX Ethernet CMA – RJ45	✓					c8-10/100eth-tx
3HE01023AA	1-port GigE CMA – SFP	✓					c1-1gb-sfp
3HE01364AA	4-port Channelized OC-3/STM-1 (DS0) ASAP MDA – SFP	✓	✓				m4-choc3-as-sfp
3HE01616AA	10-port GigE MDA – SFP Rev B [1]		✓				m10-1gb-sfp-b
3HE02021AA	1-port 10GBASE + 10-port GIGE MDA		✓				m10-1gb+1-10gb
3HE02185AA	2-port OC-3c/STM-1c/OC-12c/STM-4c CMA – SFP	✓					c2-oc12/3-sfp
3HE02499AA	1-port Channelized OC-12/STM-4 ASAP MDA	✓	✓				m1-choc12-as-sfp
3HE02500AA	12-port Channelized DS3/E3 ASAP MDA	✓	✓				m12-chds3-as
3HE02501AA	4-port Channelized DS3/E3 ASAP MDA	✓	✓				m4-chds3-as
3HE03077AA	1-port Channelized OC-3/STM-1 CES CMA	✓					c1-choc3-sfp

Table 9 **MDAs, CMAs, and ISAs Supported in 7750 SR (Continued)**

Nokia Part #	Description	SR-c4/c12	iom3-xp/-b/-c	iom4-e, iom4-e-b, iom4-e-hs, and SR-1e/2e/3e (iom-e)	SR-a4/a8 (iom-a)	SR-1	CLI String (MDA)
3HE03078AA	1-port Channelized OC-3/STM-1 CES MDA		✓				m1-choc3-ces-sfp
3HE03079AA	7750 SR 4-port CH OC-3/STM-1 CES SFP MDA	✓	✓				m4-choc3-ces-sfp
3HE03609AA	1-port GE SFP – CMA-XP	✓					c1-1gb-xp-sfp
3HE03610AA	5-port GE SFP – CMA-XP	✓					c5-1gb-xp-sfp
3HE03611AA	7750 SR 10-port GE – SFP MDA-XP	✓	✓				m10-1gb-xp-sfp
3HE03612AA	7750 SR 20-port GE – SFP MDA-XP	✓	✓				m20-1gb-xp-sfp
3HE03613AA	7750 SR 20-port GE – Copper/TX MDA-XP	✓	✓				m20-1gb-xp-tx
3HE03685AA	7750 SR 2-port 10GBASE – XFP MDA-XP	✓	✓				m2-10gb-xp-xfp
3HE03686AA	7750 SR 4-port 10GBASE – XFP MDA-XP		✓				m4-10gb-xp-xfp
3HE04272AA	7750 SR 1-port OC-12/STM-4 CES MDA	✓	✓				m1-choc12-ces-sfp
3HE04274AA	7750 SR 1-port 10GBASE – XFP MDA-XP	✓	✓				m1-10gb-xp-xfp
3HE04922AA	7750 SR / 7450 ESS Multiservice ISA [2]	✓	✓				isa-ms
3HE05142AA	7750 SR / 7450 ESS Multiservice ISA-E (no encryption) [2]	✓	✓				isa-ms-e
3HE05160AA	7750 SR 48-port 10/100/1000 – mini-RJ21 MDA-XP		✓				m48-1gb-xp-tx
3HE05942AA	7750 SR / 7450 ESS Versatile Services Module XP (VSM-CCA-XP)		✓				vsm-cca-xp
3HE05943AA	7750 SR 16-port OC-3/12c STM-1/4c POS MDA – SFP Rev B	✓	✓				m16-oc12/3-sfp-b
3HE05944AA	7750 SR 16-port ATM OC-3c/STM-1c MDA – SFP Rev B		✓				m16-atmoc3-sfp-b

Table 9 **MDAs, CMAs, and ISAs Supported in 7750 SR (Continued)**

Nokia Part #	Description	SR-c4/c12	iom3-xp/-b/-c	iom4-e, iom4-e-b, iom4-e-hs, and SR-1e/2e/3e (iom-e)	SR-a4/a8 (iom-a)	SR-1	CLI String (MDA)
3HE05945AA	7750 SR 4-port ATM OC-12c/STM-4c MDA – SFP Rev B	✓	✓				m4-atmoc12/3-sf-b
3HE05946AA	7750 SR 4-port OC-48c/STM-16c POS MDA – SFP Rev B	✓	✓				m4-oc48-sfp-b
3HE05947AA	7750 SR 2-port OC-192/STM-64 – XFP MDA-XP		✓				m2-oc192-xp-xfp
3HE06432AA	7750 SR 10-port GE SFP HS-MDAv2		✓				m10-1gb-hs-sfp-b
3HE06433AA	7750 SR 1-port 10GE HS-MDAv2		✓				m1-10gb-hs-xfp-b
3HE06521AA	2-port OC-3c/STM-1c/OC-12c/STM-4c CMA – SFP Rev B	✓					c2-oc12/3-sfp-b
3HE07282AA	7750 SR 2-port 10GE XFP + 12-port GE SFP – MDA-XP		✓				m12-1gb+2-10gb-xp
3HE07284AA	7750 SR 12-port GigE – SFP MDA-XP		✓				m12-1gb-xp-sfp
3HE08220AA	8-port Channelized DS1/E1 CMA Rev B	✓					c8-chds1
3HE09203AA	7750 SR-a 1-port 100GE MDA-a XP – CFP				✓		maxp1-100gb-cfp
3HE09204AA	7750 SR-a 10-port 10GE MDA-a XP – SFP+				✓		maxp10-10gb-sfp+
3HE09205AA	7750 SR-a 2-port 10GE SFP+ + 12-port GE SFP MDA-a				✓		ma2-10gb-sfp+12-1gb-sfp
3HE09206AA	7750 SR-a 20-port 10/100/1000 TX MDA-a – RJ45				✓		ma20-1gb-tx
3HE09207AA	7750 SR-a 22-port GE SFP/44-port GE MDA-a – CSFP				✓		ma44-1gb-csfp
3HE09240AA	7750 SR-a 4-port 10GE MDA-a – SFP+				✓		ma4-10gb-sfp+
3HE09241AA	7750 SR-a 6-port 10GE SFP+ + 1-port 40GE QSFP+ MDA-a XP				✓		maxp6-10gb-sfp+1-40gb-qsfp+

Table 9 **MDAs, CMAs, and ISAs Supported in 7750 SR (Continued)**

Nokia Part #	Description	SR-c4/c12	iom3-xp/-b/-c	iom4-e, iom4-e-b, iom4-e-hs, and SR-1e/2e/3e (iom-e)	SR-a4/a8 (iom-a)	SR-1	CLI String (MDA)
3HE09649AA	MDA-e 10-port 10 GE SFP+			✓			me10-10gb-sfp+
3HE09881AA	MDA-e 1-port 100 GE CFP2			✓			me1-100gb-cfp2
3HE10421AA	MDA–a XP - 7750 SR 1-PT 100GE CFP2				✓		maxp1-100gb-cfp2
3HE10422AA	MDA–a XP - 7750 SR 1-PT 100GE CFP4				✓		maxp1-100gb-cfp4
3HE10427AA	ISA - 7750 SR MS-ISA2 [3]			✓			me-isa2-ms
3HE10428AA	ISA - 7750 SR MS-ISA2-E [3]			✓			me-isa2-ms-e
3HE10429AA	MDA-e 6-port 10GE SFP+			✓			me6-10gb-sfp+
3HE10642AA	MDA-e 20-port GE SFP/40-port GE cSFP			✓			me40-1gb-csfp
3HE11030AA	MDA-e 2-port 100GE CFP4			✓			me2-100gb-cfp4
3HE11031AA	MDA-e 2-port 100GE QSFP28			✓			me2-100gb-qsfp28
3HE11903AA	MDA-e 12-port 10/1GE MACsec SFP+			✓			me12-10/1gb-sfp+
3HE12333AA	**MDA-e-XP - 7750 SR 6-PT 100GE QSFP28**					✓	**me6-100gb-qsfp28**
3HE12334AA	**MDA-e-XP - 7750 SR 12-PT 100GE QSFP28**					✓	**me12-100gb-qsfp28**

Notes:

1. This card is support-discontinued, but still compatible with SR OS.
2. See Usage Notes for specifics.
3. Only Ethernet MDA-e cards are supported in IOM4-e-hs. ISAs are not supported.

Table 10 **MDAs, ISAs, and VSMs Supported in 7450 ESS in Non-Mixed Mode**

Nokia Part #	Description	iom3-xp/-b/-c	iom4-e and iom4-e-b	CLI String (MDA)
3HE00021AA	7750 SR 60-port 10/100TX MDA – mini-RJ21 [1]	✓		m60-10/100eth-tx
3HE00023AA	7750 SR 20-port 100FX MDA – SFP [1]	✓		m20-100eth-sfp
3HE00030AA	7750 SR 1-port 10GBASE-LW/LR MDA with optics – Simplex SC [1]	✓		m1-10gb
3HE00033AA	7750 SR 16-port OC-3c/STM-1c MDA – SFP [1]	✓		m16-oc3-sfp
3HE00037AA	7750 SR 8-port OC-12c/STM-4c MDA – SFP [1]	✓		m8-oc12/3-sfp
3HE00038AA	7750 SR 16-port OC-12c/STM-4c MDA – SFP [1]	✓		m16-oc12/3-sfp
3HE00043AA	7750 SR 2-port OC-48c/STM-16c MDA – SFP [1]	✓		m2-oc48-sfp
3HE00044AA	7750 SR 4-port OC-48c/STM-16c MDA – SFP [1]	✓		m4-oc48-sfp
3HE00048AA	7750 SR 1-port OC-192c/STM-64c MDA with SR-1/I-64.1 optic – Simplex SC [1]	✓		m1-oc192
3HE00049AA	7750 SR 1-port OC-192c/STM-64c MDA with IR-2/S-64.2 optic – Simplex SC [1]	✓		m1-oc192
3HE00230AA	60-port 10/100TX MDA – mini-RJ21	✓		m60-10/100eth-tx
3HE00231AA	20-port 100FX MDA – SFP	✓		m20-100eth-sfp
3HE00235AA	1-port 10GBASE-LW/LR MDA with optics – Simplex SC [2]	✓		m1-10gb
3HE00237AA	16-port OC-3c/STM-1c MDA – SFP	✓		m16-oc3-sfp
3HE00238AA	8-port OC-12c/STM-4c MDA – SFP	✓		m8-oc12/3-sfp
3HE00239AA	2-port OC-48c/STM-16c MDA – SFP	✓		m2-oc48-sfp
3HE00243AA	16-port OC-12c/STM-4c MDA – SFP	✓		m16-oc12/3-sfp
3HE00244AA	4-port OC-48c/STM-16c MDA – SFP	✓		m4-oc48-sfp
3HE00709AA	7750 SR 1-port OC-192c/STM-64c MDA with LR-2/L-64.2 optic – Simplex SC [1]	✓		m1-oc192

Table 10 **MDAs, ISAs, and VSMs Supported in 7450 ESS in Non-Mixed Mode (Continued)**

Nokia Part #	Description	iom3-xp/-b/-c	iom4-e and iom4-e-b	CLI String (MDA)
3HE01532AA	10-port GigE MDA – SFP Rev B [2]	✓		m10-1gb-sfp-b
3HE01616AA	7750 SR 10-port GigE MDA – SFP Rev B [1]	✓		m10-1gb-sfp-b
3HE02021AA	7750 SR 1-port 10GBASE + 10-port GIGE MDA [1]	✓		m10-1gb+1-10gb
3HE02022AA	7450 ESS 1-port 10GBASE+10-port GigE MDA	✓		m10-1gb+1-10gb
3HE03611AA	7750 SR 10-port GE – SFP MDA-XP [1]	✓		m10-1gb-xp-sfp
3HE03612AA	7750 SR 20-port GE – SFP MDA-XP [1]	✓		m20-1gb-xp-sfp
3HE03613AA	7750 SR 20-port GE – Copper/TX MDA-XP [1]	✓		m20-1gb-xp-tx
3HE03614AA	7450 ESS 10-port GE – SFP MDA-XP	✓		m10-1gb-xp-sfp
3HE03615AA	7450 ESS 20-port GE – SFP MDA-XP	✓		m20-1gb-xp-sfp
3HE03616AA	7450 ESS 20-port GE – Copper/TX MDA-XP	✓		m20-1gb-xp-tx
3HE03685AA	7750 SR 2-port 10GBASE – XFP MDA-XP [1]	✓		m2-10gb-xp-xfp
3HE03686AA	7750 SR 4-port 10GBASE – XFP MDA-XP [1]	✓		m4-10gb-xp-xfp
3HE03687AA	7450 ESS 2-port 10GBASE – XFP MDA-XP	✓		m2-10gb-xp-xfp
3HE03688AA	7450 ESS 4-port 10GBASE – XFP MDA-XP	✓		m4-10gb-xp-xfp
3HE04273AA	7450 1-port 10GBASE – XFP MDA-XP	✓		m1-10gb-xp-xfp
3HE04274AA	7750 SR 1-port 10GBASE – XFP MDA-XP [1]	✓		m1-10gb-xp-xfp
3HE04922AA	7750 SR / 7450 ESS Multiservice ISA [3]	✓		isa-ms
3HE05142AA	7750 SR / 7450 ESS Multiservice ISA-E (no encryption) [3]	✓		isa-ms-e
3HE05159AA	7450 SR 48-port 10/100/1000 – mini-RJ21 – MDA-XP	✓		m48-1gb-xp-tx
3HE05160AA	7750 SR 48-port 10/100/1000 – mini-RJ21 MDA-XP [1]	✓		m48-1gb-xp-tx
3HE05942AA	7750 SR / 7450 ESS Versatile Services Module XP (VSM-CCA-XP)	✓		vsm-cca-xp

Table 10 **MDAs, ISAs, and VSMs Supported in 7450 ESS in Non-Mixed Mode (Continued)**

Nokia Part #	Description	iom3-xp/-b/-c	iom4-e and iom4-e-b	CLI String (MDA)
3HE05943AA	7750 SR 16-port OC-3/12c STM-1/4c POS MDA – SFP Rev B [1]	✓		m16-oc12/3-sfp-b
3HE05946AA	7750 SR 4-port OC-48c/STM-16c POS MDA – SFP Rev B [1]	✓		m4-oc48-sfp-b
3HE06382AA	7450 ESS 16-port OC-3/12c STM-1/4c POS MDA – SFP Rev B	✓		m16-oc12/3-sfp-b
3HE06383AA	7450 ESS 4-port OC-48c/STM-16c POS MDA – SFP Rev B	✓		m4-oc48-sfp-b
3HE06432AA	7750 SR 10-port GE SFP HS-MDAv2 [1]	✓		m10-1gb-hs-sfp-b
3HE06434AA	7450 ESS 10-port GE SFP HS-MDAv2	✓		m10-1gb-hs-sfp-b
3HE06435AA	7450 ESS 1-port 10GE HS-MDAv2	✓		m1-10gb-hs-sfp-b
3HE07282AA	7750 SR 2-port 10GE XFP + 12-port GE SFP – MDA-XP [1]	✓		m12-1gb+2-10gb-xp
3HE07283AA	7450 ESS 2-port 10GE XFP + 12-port GE SFP – MDA-XP	✓		m12-1gb+2-10gb-xp
3HE07284AA	7750 SR 12-port GigE – SFP MDA-XP [1]	✓		m12-1gb-xp-sfp
3HE07285AA	7450 ESS 12-port GigE – SFP MDA-XP	✓		m12-1gb-xp-sfp
3HE09649AA	MDA-e 10-port 10 GE SFP+		✓	me10-10gb-sfp+
3HE09881AA	MDA-e 1-port 100 GE CFP2		✓	me1-100gb-cfp2
3HE10427AA	ISA - 7750 SR MS-ISA2		✓	me-isa2-ms
3HE10428AA	ISA - 7750 SR MS-ISA2-E		✓	me-isa2-ms-e
3HE10429AA	MDA-e 6-port 10GE SFP+		✓	me6-10gb-sfp+
3HE10642AA	MDA-e 20-port GE SFP/40-port GE cSFP		✓	me40-1gb-csfp
3HE11030AA	MDA-e 2-port 100GE CFP4		✓	me2-100gb-cfp4
3HE11031AA	MDA-e 2-port 100GE QSFP28		✓	me2-100gb-qsfp28
3HE11903AA	MDA-e 12-port 10/1GE MACsec SFP+		✓	me12-10/1gb-sfp+

Notes:

1. Supported only with 7750 SR IOM3-XP in the 7450 ESS chassis.
2. This card is support-discontinued, but still compatible with SR OS.
3. See Usage Notes for specifics.

The following table summarizes the MDAs, ISAs, and VSMs supported in SR OS for the 7450 ESS in mixed mode. 7750 SR MDAs must be configured in the 7750 SR IOM3-XP, 7750 SR IOM4-e, or 7750 SR IOM4-e-B for mixed mode functionality.

Table 11 **MDAs, ISAs, and VSMs Supported in the 7450 ESS in Mixed Mode**

Nokia Part #	Description	CLI String (MDA)
3HE00021AA	60-port 10/100TX MDA – mini-RJ21	m60-10/100eth-tx
3HE00023AA	20-port 100FX MDA – SFP	m20-100eth-sfp
3HE00030AA	1-port 10GBASE-LW/LR MDA with optics – Simplex SC [1]	m1-10gb
3HE00032AA	8-port OC-3c/STM-1c MDA – SFP	m8-oc3-sfp
3HE00033AA	16-port OC-3c/STM-1c MDA – SFP	m16-oc3-sfp
3HE00037AA	8-port OC-12c/STM-4c MDA – SFP	m8-oc12/3-sfp
3HE00038AA	16-port OC-12c/STM-4c MDA – SFP	m16-oc12/3-sfp
3HE00043AA	2-port OC-48c/STM-16c MDA – SFP	m2-oc48-sfp
3HE00044AA	4-port OC-48c/STM-16c MDA – SFP	m4-oc48-sfp
3HE00048AA	1-port OC-192c/STM-64c MDA with SR-1/I-64.1 optic – Simplex SC	m1-oc192
3HE00049AA	1-port OC-192c/STM-64c MDA with IR-2/S-64.2 optic – Simplex SC	m1-oc192
3HE00071AA	4-port ATM OC-12c/STM-4c MDA – SFP	m4-atmoc12/3-sfp
3HE00074AA	16-port ATM OC-3c/STM-1c MDA – SFP	m16-atmoc3-sfp
3HE00709AA	1-port OC-192c/STM-64c MDA with LR-2/L-64.2 optic – Simplex SC [2]	m1-oc192
3HE01364AA	4-port Channelized OC-3/STM-1 (DS0) ASAP MDA – SFP	m4-choc3-as-sfp
3HE01616AA	10-port GigE MDA – SFP Rev B	m10-1gb-sfp-b
3HE02021AA	1-port 10GBASE + 10-port GIGE MDA	m10-1gb+1-10gb

Table 11 MDAs, ISAs, and VSMs Supported in the 7450 ESS in Mixed Mode (Continued)

Nokia Part #	Description	CLI String (MDA)
3HE02499AA	1-port Channelized OC-12/STM-4 ASAP MDA	m1-choc12-as-sfp
3HE02500AA	12-port Channelized DS3/E3 ASAP MDA	m12-chds3-as
3HE02501AA	4-port Channelized DS3/E3 ASAP MDA	m4-chds3-as
3HE03078AA	1-port Channelized OC-3/STM-1 CES MDA	m1-choc3-ces-sfp
3HE03079AA	7750 SR 4-port CH OC-3/STM-1 CES SFP MDA	m4-choc3-ces-sfp
3HE03611AA	7750 SR 10-port GE – SFP MDA-XP	m10-1gb-xp-sfp
3HE03612AA	7750 SR 20-port GE – SFP MDA-XP	m20-1gb-xp-sfp
3HE03613AA	7750 SR 20-port GE – Copper/TX MDA-XP	m20-1gb-xp-tx
3HE03685AA	7750 SR 2-port 10GBASE – XFP MDA-XP	m2-10gb-xp-xfp
3HE03686AA	7750 SR 4-port 10GBASE – XFP MDA-XP	m4-10gb-xp-xfp
3HE04272AA	7750 SR 1-port OC-12/STM-4 CES MDA	m1-choc12-ces-sfp
3HE04274AA	7750 SR 1-port 10GBASE – XFP MDA-XP	m1-10gb-xp-xfp
3HE04922AA	7750 SR / 7450 ESS Multiservice ISA [2, 3]	isa-ms
3HE05142AA	7750 SR / 7450 ESS Multiservice ISA-E (no encryption) [2]	isa-ms-e
3HE05160AA	7750 SR 48-port 10/100/1000 – mini-RJ21 MDA-XP	m48-1gb-xp-tx
3HE05942AA	7750 SR / 7450 ESS Versatile Services Module XP (VSM-CCA-XP)	vsm-cca-xp
3HE05943AA	7750 SR 16-port OC-3/12c STM-1/4c POS MDA – SFP Rev B	m16-oc12/3-sfp-b
3HE05944AA	7750 SR 16-port ATM OC-3c/STM-1c MDA – SFP Rev B	m16-atmoc3-sfp-b
3HE05945AA	7750 SR 4-port ATM OC-12c/STM-4c MDA – SFP Rev B	m4-atmoc12/3-sf-b
3HE05946AA	7750 SR 4-port OC-48c/STM-16c POS MDA – SFP Rev B	m4-oc48-sfp-b

Table 11 **MDAs, ISAs, and VSMs Supported in the 7450 ESS in Mixed Mode (Continued)**

Nokia Part #	Description	CLI String (MDA)
3HE05947AA	7750 SR 2-port OC-192/STM-64 – XFP MDA-XP	m2-oc192-xp-xfp
3HE06432AA	7750 SR 10-port GE SFP HS-MDAv2	m10-1gb-hs-sfp-b
3HE06433AA	7750 SR 1-port 10GE HS-MDAv2	m1-10gb-hs-xfp-b
3HE07282AA	7750 SR 2-port 10GE XFP + 12-port GE SFP – MDA-XP	m12-1gb+2-10gb-xp
3HE07284AA	7750 SR 12-port GigE – SFP MDA-XP	m12-1gb-xp-sfp
3HE09649AA	MDA-e 10-port 10GE SFP+	me10-10gb-sfp+
3HE09881AA	MDA-e 1-port 100GE CFP2	me1-100gb-cfp2
3HE10427AA	ISA - 7750 SR MS-ISA2	me-isa2-ms
3HE10428AA	ISA - 7750 SR MS-ISA2-E	me-isa2-ms-e
3HE10429AA	MDA-e 6-port 10GE SFP+	me6-10gb-sfp+
3HE10642AA	MDA-e 20-port GE SFP/40-port GE cSFP	me40-1gb-csfp
3HE11030AA	MDA-e 2-port 100GE CFP4	me2-100gb-cfp4
3HE11031AA	MDA-e 2-port 100GE QSFP28	me2-100gb-qsfp28
3HE11903AA	MDA-e 12-port 10/1GE MACSsec SFP+	me12-10/1gb-sfp+

Notes:

1. This card is support-discontinued, but still compatible with SR OS.
2. MS-ISAs and ISA applications using MS-ISAs are not supported in mixed mode with the exception of Application Assurance, IPsec, and NAT. IPsec is not supported with MS-ISA-E.
3. Starting with Release 8.0.R5, MS-ISA cards (3HE04922AA) replace IPsec-ISA cards (3HE03080AA).

2.4 Supported 7210 SAS-Sx Satellites

The following table summarizes the 7210 SAS-Sx satellites supported in SR OS.

Table 12 **7210 SAS-Sx Satellites**

Nokia Part #	Description	sat-type	Initial 7210 Software Release	Initial 7750 Host Support
3HE10328AA	SYS - 7210 SAS-Sx SONET/SDH ETR DC	ts4-choc3-sfp ts4-chstm1-sfp ts1-choc12-sfp ts1-chstm4-sfp	8.0.R4 (7705 SAR software)	15.0.R1
3HE10492AA	SYS - 7210 SAS-Sx 46F2C4SFP+	es48-1gb-sfp	8.0.R6	14.0.R4
3HE10493AA	SYS - 7210 SAS-Sx 22F2C4SFP+	es24-1gb-sfp	8.0.R6	14.0.R4
3HE10494AA	SYS - 7210 SAS-Sx 48T4SFP+	es48-1gb-tx	8.0.R9	14.0.R6
3HE10495AA	SYS - 7210 SAS-Sx 24T4SFP+	es24-1gb-tx	8.0.R9	14.0.R6
3HE10496AA	SYS - 7210 SAS-Sx 48Tp4SFP+ (PoE)	es48-1gb-tx	8.0.R8	14.0.R6
3HE10497AA	SYS - 7210 SAS-Sx 24Tp4SFP+ (PoE)	es24-1gb-tx	8.0.R8	14.0.R6
3HE10835AA	SYS - 7210 SAS-Sx 64SFP+ 4CFP4	es64-10gb-sfpp+4-100gb-cfp4	9.0.R7	15.0.R4
3HE10530AA	SYS - 7210 SAS-S 48F4SFP+ (AC)	es48-sass-1gb-sfp	9.0.R7	15.0.R4
3HE10531AA	SYS - 7210 SAS-S 48F4SFP+ (DC -48)	es48-sass-1gb-sfp	9.0.R7	15.0.R4
3HE10532AA	SYS - 7210 SAS-S 24F4SFP+ (AC)	es24-sass-1gb-sfp	9.0.R7	15.0.R4
3HE10533AA	SYS - 7210 SAS-S 24F4SFP+ (DC -48)	es24-sass-1gb-sfp	9.0.R7	15.0.R4
3HE10534AA	SYS - 7210 SAS-S 48T4SFP+ AC	es48-1gb-tx	9.0.R7	15.0.R4
3HE10535AA	SYS - 7210 SAS-S 48T4SFP+ -48VDC	es48-1gb-tx	9.0.R7	15.0.R4
3HE10536AA	SYS - 7210 SAS-S 24T4SFP+ AC	es24-1gb-tx	9.0.R7	15.0.R4
3HE10537AA	SYS - 7210 SAS-S 24T4SFP+ -48VDC	es24-1gb-tx	9.0.R7	15.0.R4

3 New Features

The following sections describe the new features added in SR OS releases. New Features from Releases 15.0.R1 to 15.0.R6 also apply to Release 15.1. Refer to the most recent *SR OS 15.0 Release Notes* for the summary of new features in Releases 15.0.R1 through 15.0.R6.

Note: New features that were added in earlier releases, but which were not documented until the current release, are marked **[NEW]** and are documented in the section for the applicable release.

3.1 Release 15.1.R3

Release 15.1.R3 has no new major features. See also Enhancements in Release 15.1.R3 and Resolved Issues in Release 15.1.R3.

3.2 Release 15.1.R2

3.2.1 Hardware

3.2.1.1 MDA-e-XP 12-port 100GE QSFP28

Release 15.1.R2 introduces the MDA-e-XP 12-port 100GE QSFP28. This MDA is supported on the 7750 SR-1 chassis and supports a wide range of pluggable interfaces. In Release 15.1.R2, the 1x100G and 4x10G pluggable interfaces are supported.

3.2.1.2 APEQ-DC-4275, Quad PCM, and PCM Fan tray

Release 15.1.R2 introduces the LVDC Advanced Power Equalizer Module (APEQ) APEQ-DC-4275 for the 7950 XRS-20e system. This APEQ, along with the supporting Quad Power Connect Module (Quad PCM) and the PCM Fan hardware, offers a complete redundant powering solution for a 7950 XRS-20e system. The 7950 XRS-20e universal chassis can be deployed with up to twelve APEQs and twelve Quad PCMs. Two PCM Fans must always be installed when deploying APEQs and Quad PCMs (to offer 1+1 redundancy).

This feature was incorrectly included as new in Release 15.1.R1-1. Release 15.1.R2 includes the full support.

3.3 Release 15.1.R1-1

3.3.1 Hardware

3.3.1.1 7750 SR-1

Release 15.1.R1-1 introduces the support for 7750 SR-1, a new chassis in the 7750 SR platform family. The 7750 SR-1 has an integrated CPM with IOM, supports up to two MDA-e-XP cards, and enables 2:1 capacity oversubscription with intelligent fan-in/fan-out capabilities.

The chassis provides full synchronization with both SyncE and IEEE 1588 as well as Nokia 7210 Service Access Switch-S/-Sx (SAS-S/-Sx) satellite system support. 7750 SR-1, equipped with modular rear-mounted fans, has two different assemblies, depending on the power source:

- The AC variant has two rear mounted modular power supplies
- The DC variant comes with integrated dual feeds at the rear of the system

7750 SR-1 supports 128k ingress and 128k egress queues shared across two MDA-e-XP cards.

7750 SR-1 is offered as:

Table 13 **7750 SR-1 Chassis**

Nokia Part #	Description
3HE12298AA	SYS - 7750 SR-1 DC - Core Router 1.2TB
3HE12298BA	SYS - 7750 SR-1 DC - Edge Router 1.2TB
3HE12298CA	SYS - 7750 SR-1 DC High Scale Edge 1.2TB
3HE12299AA	SYS - 7750 SR-1 AC - Core Router 1.2TB
3HE12299BA	SYS - 7750 SR-1 AC - Edge Router 1.2TB
3HE12299CA	SYS - 7750 SR-1 AC High Scale Edge 1.2TB

3.3.1.2 MDA-e-XP 6-port 100GE QSFP28

Release 15.1.R1-1 introduces the MDA-e-XP 6-port 100GE QSFP28. This MDA is supported on the 7750 SR-1 chassis and supports a wide range of pluggable interfaces. In Release 15.1.R1, the 1x100G and 4x10G pluggable interfaces are supported.

3.3.1.3 QSFP28 Breakout Connectors

Release 15.1.R1-1 adds the support for assemblies that use QSFP28 connectors. In order to use a port on these assemblies, the connector type must first be configured. Connectors come in single-port and multiple-port breakout-connectors varieties. Ports belonging to a connector can be used like other physical Ethernet ports within SR OS. In Release 15.1.R1-1, the 1x100G and 4x10G connector types are supported.

3.3.2 System

3.3.2.1 CLI Log Destination

Release 15.1.R1-1 introduces CLI log capability that outputs log events to a CLI session. An operator can subscribe to a CLI log from within a CLI session using the **tools perform log subscribe-to log-id** *log-id* command. The log events are sent to the CLI session for the duration of that CLI session (or until an **unsubscribe-from** command is issued).

3.3.2.2 Filter Policies: TCP Flags

Release 15.1.R1-1 introduces the support for the following additional TCP flag match criteria (as defined in RFC 793/3168/3540) for IPv4 and IPv6 ACL filter policies: **tcp-fin**, **tcp-rst**, **tcp-psh**, **tcp-urg**, **tcp-ece**, **tcp-cwr**, and **tcp-ns**. These new match criteria are supported on FP4-based line cards only.

3.3.2.3 NETCONF Notifications

Release 15.1.R1-1 introduces the ability for a NETCONF client to receive asynchronous notifications through NETCONF (as per RFC 5277) from SR OS routers in a structured XML format based on YANG data models. A NETCONF client needs to maintain an open NETCONF session with the SR OS router and send a <create-subscription> operation before notifications can be received. A pre-configured or user-configured stream can be used in the <create-subscription> operation. The following NETCONF notifications are supported:

- sros-config-change-event: sent with every configuration change (any new, deleted, or modified configuration)
- sros-state-change-event: sent with every state change
- sros-command-accounting-event: sent to keep track of which user issued which command on the SR OS router
- sros-log-generic-event: contains the rest of the SR OS log events (except for Lawful Intercept events)

3.3.2.4 Telemetry

Release 15.1.R1-1 adds the support for the latest version of gnmi.proto (version 0.4.0). In addition, the gRPC server implementation now supports multiple gRPC users within one TCP connection.

3.3.2.5 System Profiles

Release 15.1.R1-1 allows for multiple system profiles where each profile supports different capabilities for FP4-based line cards. The system profile is defined in the BOF and is used by the system after the next node reset. In Release 15.1.R1-1, the BOF **system-profile** parameter should be configured to **profile-a** on a 7750 SR-1. All other systems must use the existing system capabilities which is indicated by the omission of the **system-profile** parameter in the BOF. Contact your Nokia representative for system profile information.

4 Enhancements

The following sections describe new enhancements in SR OS releases. Enhancements from Releases 15.0.R1 to 15.0.R6 also apply to Release 15.1. Refer to the most recent *SR OS 15.0 Release Notes* for the summary of enhancements in Releases 15.0.R1 through 15.0.R6.

Note:

- For the list of new and updated Application Assurance protocols and applications supported in Release 15.1.R3, see the following spreadsheet at the Nokia online customer support site:

  SR OS 15.1 AA Protocols and Applications

  The spreadsheet may also be updated between maintenance releases to reflect recent AA protocol and application updates. To subscribe to document and spreadsheet notifications, see the online customer support site.

  For a complete list of all AA protocols and applications, contact your regional support organization.
- Enhancements that were added in earlier releases, but which were not documented until the current release, are marked **[NEW]** and are documented in the section for the applicable release.

4.1 Release 15.1.R3

4.1.1 Services General

- MACsec LAN mode is now ready for production networks. This feature was introduced in Release 15.0.R5.

4.2 Release 15.1.R2

4.2.1 Satellites

- In Release 15.1.R2, the 7210 SAS-SX SONET/SDH satellite adds the support for MC-APS (Multi-Chassis Automatic Protection Switching) with dual hosts. This mode provides protection from failures in the line (OC3/OC12/STM1/STM4), optical transceivers, TDM satellite, host, and satellite-to-host connectivity, therefore allowing higher availability for the satellite's services. [274503]

4.2.2 BGP

- Release 15.1.R2 introduces a new "unchanged" value for the Label Type field and a new Lbl Allocation field in the RIB Out section output of the show router bgp routes hunt command. [271892]

4.2.3 QoS

- Release 15.1.R2 enhances the burst control group (BCG) **show** output to include the BCG information relating to internal queues. The internal queues are included in the Recalc List Information summary and are displayed as a single entry (“TiMOS Internal Queues”) when the **member-queues** parameter is specified. [275908]

4.2.4 WLAN-GW

- In data-triggered mode with **authenticate-on-dhcp** enabled, the host running Release 15.1.R2 can switch to DHCP-triggered state when receiving control-plane traffic under following conditions. [186190]
  - when no RADIUS Access-Accept or Reject message is received, control-plane authentication is allowed after RADIUS timeout (timestamp of initial Request + **retry * timeout**)

  - when RADIUS Access-Accept or Reject message is received but promotion to DSM or ESM fails, control-plane authentication will be allowed immediately

4.3 Release 15.1.R1-1

4.3.1 IP/RTM

- Release 15.1.R1-1 adds a new **description** CLI command under the **show router interface** context to display the list of configured interfaces with the corresponding port or SAP delimiter, Administrative and Operational State, and the Interface Description string. [267414]

4.3.2 QoS

- Release 15.1.R1-1 adds the support for four additional burst control groups on FP4-based line cards with target visitation times of 50µs, 10µs, 5µs, and 1µs for both ingress and egress queues. When displaying the burst-control-group member queues on FP4-based line cards, the expected utilization bandwidth (**exp-util-bw**) parameter is not applicable.

4.3.3 Routing

- NG-MVPN Core Diversity is now ready for production networks. This feature was introduced in Release 15.0.R4.

4.3.4 Subscriber Management

- Release 15.1.R1-1 introduces a new Web Portal Protocol (WPP) port attribute. This is an ASCII string that contains up to 35 characters, including WPP system name and SAP information. This new format can be configured in the **config**>**service**>**vprn**> **wpp**>**portals**>**portal**>**port-format vendor-specific** CLI context. [265322]

- ESM over GTP is now ready for production networks. This feature was introduced in Release 15.0.R4.

4.3.5 Services

- Limited Operation State for IPsec Gateway or IPsec Tunnel is now ready for production networks. This feature was introduced in Release 15.0.R4.

4.3.6 MPLS

- In Release 15.1.R1-1, an interface in an MPLS or RSVP context remains in a down state if the user enters a **no address** or **no unnumbered** command in the **config**>**router**>**interface** CLI context. In prior releases, the interface was incorrectly deleted from the MPLS and RSVP contexts. [214303]
- Release 15.1.R1-1 extends the length of the MPLS **lsp** *lsp-name* parameter from 32 characters to 64 characters.

4.3.7 Scaling

The following scaling numbers have been increased; contact your Nokia representative for details:

- The maximum number of ACL filter entries per system
- The maximum number of event logs per system

5 Limited Support Features

This section describes the SR OS features that are intended for laboratory use only and which should not be used in production networks.

See also Unsupported Features and Known Limitations for more information about features that may not be fully supported.

Table 14 **Limited Support Features**

Section	Feature	Release Introduced
Hardware	virtualized Simulator (vSim)	12.0.R4
System	NETCONF candidate datastore support (Transactions) and the Nokia SR OS YANG data model	14.0.R1
	IEEE 1588 on 7750 SR-1	15.1.R1-1
Services	Perfect Stream	8.0.R4
	Ad Insertion (ADI)	8.0.R4
	MACsec XPN mode	15.0.R5
Subscriber Management	PPPoE Client for vRGW	15.0.R4
	Home LAN Extensions Note: Limited Support is only applicable to the EVPN portion	15.0.R4
Application Assurance	DEM WLAN-GW Access Network Location for Policy and Reporting	15.0.R4

6 Unsupported Features

The following tables summarize the features that are not supported on certain SR OS platforms (marked by an X where unsupported). All SR OS features are supported on all platforms unless otherwise listed in the table below.

Some platforms do not support applications using ISAs; see also Release 15.1.R3 Supported Hardware and Usage Notes for more information.

6.1 Hardware

Table 15 **Unsupported Hardware Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
ATM interfaces, MDA, and services	X			X		X	X		X
ASAP MDAs and associated interface types	X			X		X	X		X
CES MDAs and associated interface types	X			X		X	X		X
Channelized and TDM interfaces	X			X		X	X		X
VSM Cross-Connect Aggregation (CCA)	X			X	X	X		X [1]	X

Note:

1. When mixed mode is enabled, VSM is only supported using the 7750 SR VSM-CCA-XP module in a 7750 SR IOM.

6.2 System

Table 16 **Unsupported System Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
IEEE 1588 PTP	X[1]	[3]			X[2]		[3]	[3]	
IEEE 1588 Port-Based Timestamping (PBT)	X[4]				X				
ACL Filter Egress Rate-Limit Action				X					
BITS input port redundancy					X[6]				X[5]
BITS Out support					X[6]				
Centralized (CPM-based) CPU-Protection.				X[7]	X[7]	X[7]			X[7]
Forwarding Path Extension (FPE) IP GRE Tunnel without ISA Port Cross-Connect (PXC)[8], [16]	[9]	[10]	[10]	[11]	X	[12]	X	[10], [13]	
Hybrid OpenFlow Switch							X		
Ingress Multicast Path Management				X	X	X			X
OOB Management Ethernet Port Redundancy				X	X	X			X
Ethernet Satellites[14]		[16]	[16]		X		X		
SONET/SDH Satellites[15]	X		X				X	X	X
Soft Reset					X				X
Sub-second CCM-enabled MEPs					X				
System Alarm Contact Inputs	X	X	X		X	X	X	X	X

Notes:

1. Not supported on the 7950 XRS-16c; supported on the 7950 XRS-20/20e and XRS-40.
2. Not supported on 7750 SR-c12 with CFM-XP; supported on the 7750 SR-c4 and on the 7750 SRc-12 with CFM-XP-B.
3. Support for IEEE 1588 requires at least a CPM3 or later. The CPM3 must also include PCN C04764. Earlier CPMs do not support IEEE 1588.

   For more information about the PCN, refer to https://services.support.alcatel-lucent.com/services/pcn/PCN-001/cgi-bin/pdfpcn.cgi?C04764.pdf+C04764
4. Not supported on the 7950 XRS-16c or on the extension chassis of the 7950 XRS-40; supported on the 7950 XRS-20/20e and the master chassis of the 7950 XRS-40.
5. BITS Input is supported, but there is no redundancy on a 7750 SR-1.
6. Supported on the 7750 SR-c4 but not supported on the 7750 SR-c12.
7. Note that Distributed CPU Protection (DCP) is supported.
8. PXC is supported as follows:
   - only on Ethernet ports
   - not on ports in DWDM, WAN, or OTN mode
9. Requires 10G and 100G ports.
10. Requires SF/CPM5, 10G ports and 100G ports on FP3-based cards.
11. Requires 10G ports, excluding the ports on the following MDA: 7750 SR 4-PT 10GE SFP+ MDA-a (ma4-10gb-sfp+).
12. PXC requires 10G ports and above.
13. Requires SR line cards configured as 7450 mixed-mode.
14. 7210 Ethernet satellites use 7210 SAS Release 8.0 or 9.0. The 10GE Ethernet satellite requires 7210 SAS Release 9.0.R4 or higher.
15. 7210 SONET and SDH satellites use 7705 SAR Release 8.0.R4 or higher.
16. Not supported on IOM4-e-HS.

6.3 Quality of Service

Table 17 **Unsupported QoS Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
Named Pools	X	[1]	[1]	X	X	X			X
Ingress shared queuing (Dual-Pass)	X	[1]	X [1, 2]						X
Policers (except for Distributed CPU Protection)				X					

Notes:

1. Not supported on IOM4-e-HS.
2. Not supported on 400G FP3 line cards.

6.4 Routing

Table 18 **Unsupported Routing Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
BGP for forwarding unicast packets in GRT							X		
BGP RFC 3107-labeled routes for forwarding unicast packet in GRT [1]							X		
ABR/RR capability for BGP RFC 3107-labeled routes [2]							X		
Cflowd							X		
IPv6 routing (unicast and multicast, 6PE, 6VPE, QoS criteria matching within a VPLS or Epipe service)							X		
IP Multicast routing and forwarding • Protocols: IGMP, MLD, PIM, and MSDP • MVPN • P2MP LSP for forwarding multicast packet in GRT and in MVPN [3]							X		

Notes:

1. BGP RFC 3107-labeled routes are supported in L2 services only.
2. LDP-BGP stitching is supported.
3. P2MP LSP is supported in VPLS.

6.5 MPLS

Table 19 **Unsupported MPLS Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
GMPLS UNI		[1]	[1]		X		X	X	
LDP IPv6							X		
P2MP LDP FEC							X [2]		
P2MP RSVP-TE LSP							X [2]		

Notes:

1. Not supported on IOM4-e-HS.
2. Only supported when used as an I-PMSI in a VPLS context.

6.6 Services

Table 20 **Unsupported Services Features**

Feature	**7950 XRS**	**7750 SR-7/12**	**7750 SR-12e**	**7750 SR-a4/a8**	**7750 SR-c4/c12**	**7750 SR-1e/2e/3e**	**7450 ESS without mixed mode**	**7450 ESS with mixed mode**	**7750 SR-1**
Circuit Emulation services (for example, Cpipe SAPs)	X			X		X			X
new-qinq-untagged-sap configurability for :*.0 and :0.0 SAPs	X [1]								
EVPN features							X		
FCC/RET/VQM	X			X	X	X	X	X	X
Full VPRN support							X		
Frame Relay interfaces and services (for example, Fpipe SAPs)	X			X		X	X		X
IP Mirroring							X		
config mirror mirror-source				X	X	X			
Tunnel services (IPsec, GRE, IP-in-IP tunnel termination) [2, 3]	X			X	X [4]		X		X
IPv6 tunnel services (IPsec, GRE, IP-in-IP tunnel termination) [2, 3]	X			X	X [4]		X		X
G.8031 (Ethernet tunnel support)		[7]	[7]		X				
Multi-Chassis features using IPsec (MC-IPsec) [5]	X			X	X [4]		X		X
sFlow				X	X	X	X	X	
Spoke termination on L3 (IES/VPRN) interfaces							X		
ECMP for VXLAN IPv4 Tunnels of R-VPLS [6]					X		X		

Notes:

1. This feature is always “on” for the 7950 XRS.
2. Requires an MS-ISA/MS-ISA2/MS-ISM.
3. Requires an isa-tunnel/isa2-tunnel application license.
4. Not supported on 7750 SR-c4 only.
5. Requires an MS-ISA/MS-ISA2/MS-ISM.
6. All of the cards in the system must be FP3-based.
7. Not supported on IOM4-e-HS.

6.7 Subscriber Management

Table 21 **Unsupported Subscriber Management Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
IPv4 local DHCP Server	X								
IPv6 local DHCP Server	X						X		
GTP Uplink	X			X	X		X	X	
L2TP LNS [1, 2]	X			X	X		X		X
port-policy command [1, 2]	X			X					X
NAT [1, 2]	X			X	X [3]		X		X
Subscriber Accumulated Statistics [4]	X				X		X	X	
Subscriber Management—Routed CO (VPRN/IES subscriber interfaces)	X						X		
Subscriber Management—Bridged CO (VPLS)	X								

Table 21 **Unsupported Subscriber Management Features (Continued)**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
vRGW on regular group interfaces [1, 2]	X			X	X		X	X	X
vRGW on WLAN-GW group interfaces [1, 2]	X			X	X		X	X	X
WLAN gateway (WLAN-GW) [1, 2]	X			X	X		X	X	X
GTP Access [5]	X			X	X		X	X	
Bonding [5]	X			X	X		X	X	

Notes:

1. Requires an MS-ISA/MS-ISA2/MS-ISM (along with -E variants on the 7750 SR).
2. Requires an isa-bb/isa2-bb application license.
3. Supported on the 7750 SR-c12 but not supported on the 7750 SR-c4.
4. Requires CPM5.
5. Requires FPE. See Table 16.

6.8 Application Assurance

Table 22 Unsupported AA Features

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-c4/c12	7750 SR-1e/2e/3e	7450 ESS without mixed mode	7450 ESS with mixed mode	7750 SR-1
Application Assurance [1]	X			X					X
AARP	X			X	X [2]				

Notes:

1. Requires an MS-ISA/MS-ISA2/MS-ISM (along with -E variants on the 7750 SR) and an isa-aa/isa2-aa application license.
2. Supported on the 7750 SR-c12 but not supported on the 7750 SR-c4.

7 Deprecated Features

The following sections describe deprecated features in SR OS releases. Deprecated features from Releases 15.0.R1 to 15.0.R6 also apply to Release 15.1. Refer to the most recent *SR OS 15.0 Release Notes* for the summary of deprecated features in Releases 15.0.R1 through 15.0.R6.

Note:

- The release image should not be loaded onto deprecated platforms.
- Deprecated hardware should be removed from the router before upgrading.
- Deprecated features should be deconfigured on the router before upgrading.

7.1 Release 15.1.R3

No features have been deprecated in Release 15.1.R3 since Release 15.1.R2.

7.2 Release 15.1.R2

No features have been deprecated in Release 15.1.R2 since Release 15.1.R1-1.

7.3 Release 15.1.R1-1

No features have been deprecated in Release 15.1.R1-1.

8 Changed or Deprecated Commands

This section describes the SR OS commands that have been changed or deprecated. Unless otherwise noted, all changed CLI commands are accepted during boot and exec and are converted on upgrade.

See also Software Upgrade Procedures for more information about the behavior of commands or parameters that have been modified or deprecated between releases. Changed or deprecated commands from Releases 15.0.R1 to 15.0.R6 also apply to Release 15.1. Refer to the most recent *SR OS 15.0 Release Notes* for the summary of changed or deprecated commands in Releases 15.0.R1 through 15.0.R6.

8.1 Release 15.1.R3

No commands have been changed or deprecated in Release 15.1.R3 since Release 15.1.R2.

8.2 Release 15.1.R2

8.2.1 Services General

In Release 15.1.R2, the **def-mesh-vc-id** parameter in the **configure service vpls** CLI command context has been removed.

8.3 Release 15.1.R1-1

8.3.1 CLI

In Release 15.1.R1-1, the **configure system management cli configuration immediate** CLI command has been deprecated. It is replaced with the **configure system management-interface cli classic-cli allow-immediate** CLI command.

8.3.2 NETCONF

In Release 15.1.R1-1, the **yang-modules** CLI command has been removed from the **configure system netconf** context. It has been added to the **configure system management-interface** context.

8.3.3 QoS

In Release 15.1.R1-1, the **policy-name** command has been removed from the **qos** contexts and replaced with the **name** parameter on the **create** line:

The following commands using **policy-name** have been deprecated:

```
configure qos network network-policy-id create
...
        policy-name policy-name
...
exit


configure qos sap-ingress | sap-egress policy-id create
...
        policy-name name
...
exit
```

The **name** parameter in the following commands is the alternative for the deprecated **policy-name** commands:

```
configure qos network network-policy-id name name create

configure qos sap-ingress | sap-egress policy-id name name create
```

8.3.4 Filters

In Release 15.1.R1-1, the **filter-name** command has been removed from the **filter** context and replaced with the **name** parameter on the **create** line.

The following command using **filter-name** have been deprecated:

```
configure filter ip-filter | ipv6-filter | mac-filter filter-id create
...
        filter-name name
...
exit
```

The **name** parameter in the following command is the alternative for the deprecated **filter-name** command:

```
configure filter ip-filter | ipv6-filter | mac-filter filter-id name name create
```

8.3.5 Services General

In Release 15.1.R1-1, the **service-name, customer-name**, and **pw-template-name** commands have been removed from the **service** context and replaced with the **name** parameter on the **create** line:

The following commands using **service-name** have been deprecated:

```
configure service apipe | epipe | fpipe | ipipe | ies | vpls | vprn service-id
customer customer-id create
...
        service-name name
...
exit

configure mirror mirror-dest service-id type mirror-type create
...
        service-name service-name
...
exit
```

The **name** parameter in the following commands is the alternative for the deprecated **service-name** commands:

```
configure service apipe | epipe | fpipe | ipipe | ies | vpls | vprn service-id
customer customer-id name name create

configure mirror mirror-dest service-id type mirror-type name name create
```

The following command using **customer-name** has been deprecated:

```
configure service customer customer-id create
```

```
...
        customer-name name
...
exit
```

The **name** parameter in the following command is the alternative for the deprecated **customer-name** command:

```
configure service customer customer-id name name create
```

The following command using **pw-template-name** has been deprecated:

```
configure service pw-template policy-id create
...
        pw-template-name name
...
exit
```

The **name** parameter in the following command is the alternative for the deprecated **pw-template-name** command:

```
configure service pw-template policy-id name name create
```

9 Software Upgrade Procedures

The following sections contain information for upgrading to the Release 15.1.R3 software version.

- Software Upgrade Notes

  Information on upgrading the router from previous versions of SR OS software including rules for upgrading firmware and any special notes for upgrading from specific earlier versions.
- AA Signatures Upgrade Procedure

  Information on upgrading ISA to a new AA-signature load.
- Standard Software Upgrade Procedure

  Procedure for performing a standard, service-affecting upgrade including updating of firmware images.

9.1 Software Upgrade Notes

The following sections describe notes for upgrading from prior versions of SR OS to Release 15.1.R3.

Note:

- Upgrade notes that apply to earlier releases, but which were not documented until the current release, are marked **[NEW]** and are documented in the section for the applicable release.
- Automatic firmware updates may occur for CPM and IOM/IMM/XCM cards running older firmware after an SR OS upgrade. The **clear card** command or physical removal of a card must not be performed until the card is operationally up after an SR OS upgrade. This procedure also applies when subsequently adding new IOMs/IMMs/XCMs (that may have older firmware) to a chassis. An event log with “firmware upgraded” message will be issued if a firmware update had occurred for a card.

The following conventions are used in configuration files:

- Deprecated commands are not flagged as errors upon reading a configuration file with deprecated commands, but these commands will not be written to a saved configuration file.
- Modified commands are read using the old format, but they are written with the new format in a configuration file; so a configuration file saved with modified commands is not compatible with earlier releases.

- Modified parameters are supported when they are read, but the modified parameters will be converted to new minimums or maximums when saved in a configuration file.

9.1.1 Upgrading to Release 15.1.R3 or Higher

- When performing a standard software upgrade from Release 14.0.R4 or higher, and there is a user-created management router instance configured, the configuration file will fail to load. The user-created management router instance must be removed prior to the upgrade. [262212]

9.1.2 Upgrading to Release 15.1.R1-1 or Higher

- The SR OS gRPC telemetry interface in Release 15.1.R1-1 has been changed to OpenConfig gnmi.proto version 0.4.0. Prior to upgrading SR OS, clients/ collectors must be updated to account for telemetry interface changes.

9.1.3 Upgrading to Release 15.0.R6 or Higher

- When upgrading to Release 15.0.R6 from an earlier release, there is a mandatory firmware upgrade for CPM-e cards. During the software upgrade, cards that require the upgrade will automatically update their firmware when they are rebooted as part of the normal software upgrade process (ISSU or non-ISSU). The firmware upgrade will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the cards are not removed while they are programming the new firmware. [258922]

9.1.4 Upgrading to Release 15.0.R5

- When upgrading to Release 15.0.R5 from Release 15.0.R4, the configuration saved with the **admin save detail** CLI command on the 7750 SR-c4/c12 platform will fail to execute after a node reboot. This issue is only present when upgrading to Release 15.0.R5 from Release 15.0.R4. [268831]

9.1.5 Upgrading to Release 15.0.R4 or Higher

- The SR OS gRPC Telemetry interface in Release 15.0.R4 has been changed to OpenConfig gnmi.proto version 0.3.1. Prior to upgrading SR OS, clients/ collectors must be updated to account for telemetry interface changes (for example, the use of TypedValue instead of Value).
- In Release 15.0.R3 or earlier, it was possible, although invalid, to configure EVPN-MPLS with IGMP snooping together with all-active multihoming or single-active multihoming with an ESI label, and PBB-EVPN with IGMP snooping together with all-active multihoming. However, these combinations are not supported in Release 15.0.R3 or earlier. From Release 15.0.R4, these combinations are supported and will default to using data-driven IGMP-snooping state synchronization in EVPN-MPLS and PBB-EVPN service. Refer to the *Layer 2 Services and EVPN Guide* for more information.
- In Release 15.0.R3 or earlier, EVPN-MPLS with PIM-snooping for IPv4 together with single-active multihoming with an ESI label is configurable but not supported. From Release 15.0.R4 onwards, this combination is supported and will default to using data-driven PIM-snooping state synchronization in EVPN-MPLS services.Refer to *Layer 2 Services and EVPN Guide* for more information.

9.1.6 Upgrading to Release 15.0.R1 or Higher

- When upgrading from prior versions of SR OS to Release 15.0.R1 and above, if the router configuration included distributed CPU protection (DCP) policies named "_default-access-policy" or "_default-network-policy" then the upgrade may fail if the number of policer resources per line card is exceeded. Operators with such pre-existing DCP policies should either delete the associated policer configuration prior to the upgrade or ensure that the policer resources per line card won’t be exceeded when applied to all access/network interfaces in the system.
- The chassis mode command was required to differentiate services and scaling available on early IOMs. As of Release 15.0, those early IOMs are no longer supported, and there is no requirement for the differentiation using the chassis-mode command. When upgrading to Release 15.0 or higher, the chassis mode shall be changed to chassis mode D, and it cannot afterwards be changed.
- The enhanced Diameter Gy Extended Failure Handling (EFH) triggers are automatically activated when EFH is enabled; therefore, EFH should be disabled prior to upgrade to Release 15.0.R1 in production networks.

- If upgrading from Release 13.0 or earlier, Nokia highly recommends that the secondary RADIUS accounting policy should copy all configuration except the server configuration of the primary accounting policy to preserve the same accounting behavior.
- With the support for 7210 SR OS Release 9.0 running Ethernet satellites, it is strongly recommended to update the firmware when upgrading Ethernet satellites to 7210 Release 9.0. This firmware update can be accomplished by performing the following steps during the upgrade procedure.
    1. Store the new 7210 SAS-S/SX image files in a new software repository location.
    2. Modify the satellite configuration to reference the new software repository.
    3. Ensure that the images are synchronized with the satellite using the **admin satellite eth-sat** *sat-id* **sync-boot-env** command.
    4. Reboot the satellite when desired using the **admin satellite eth-sat** *sat-id* **reboot upgrade** command.

       The satellite will take longer to reboot due to the firmware upgrade process.

  The general procedure for upgrading an Ethernet satellite software can be found in the *7450 ESS, 7750 SR, and 7950 XRS Basic System Configuration Guide* in the “Satellite Software Upgrade Overview” section.

9.1.7 Upgrading to Release 14.0.R7 or Higher

- When performing a standard software upgrade from Release 14.0.R4, 14.0.R5, or 14.0.R6 and there is an Optical Extension Shelf (OES) configuration, the configuration file will fail to load. These deprecated OES commands must be removed prior to an upgrade.

9.1.8 Upgrading to Release 14.0.R6 or Higher

- When upgrading to Release 14.0.R6 or higher from a previous release, there is a firmware upgrade for the XCM cards of a 7950 XRS to support IEEE 1588 PBT on XMAs and C-XMAs located in the XCM. This upgrade is not applied during the Soft Reset of an ISSU operation. A hard reset of the card is required to upgrade the firmware and enable the feature.

  The status of the firmware can be checked for each card by using the **show card detail** command and checking the Firmware revision status field, which will display “Upgrade on next hard reset” if the XCM is in this state.

  The firmware update will cause a longer reboot time than usual. [234752]

Caution: Do not remove the cards while their firmware is being upgraded.

9.1.9 Upgrading to Release 14.0.R4 or Higher

- When a router is upgraded from a release prior to Release 14.0.R4 to Release 14.0.R4 or higher, the following immediate changes can be expected.
  - On every router using BGP for IP routing (whether or not it has labeled-unicast sessions), BGP memory usage may increase slightly, proportional to the number of active IP routes that are not BGP. This increase is because the default policy of adding every active IP route (non-BGP) to the BGP RIB involves two RIBs rather than one.
  - Unlabeled routes are no longer advertised over labeled-unicast sessions and vice versa. This behavior can be restored, after upgrade, using **route-table-import** policies applied to the labeled-unicast and/or unlabeled RIBs.
  - The route or routes that are advertised to a peer for a given IP prefix may be different after the upgrade because the labeled-unicast RIB (used to find the path to advertise to labeled-unicast peers) generally does not have a view of all paths for that prefix in the unlabeled RIB. The converse is also true. The unlabeled RIB (used to find the path to advertise to unlabeled peers) generally does not have a view of all paths for that prefix in the labeled-unicast RIB.
  - It is no longer a mandatory requirement that a BGP route for an IP destination be used in the route table (for IP forwarding) for it to be advertised. If a BGP route is the best path and the used route for the IP prefix has not been imported from the route table, then the best BGP route in the BGP RIB is advertisable to peers without any further configuration.
  - Received BGP routes with any AFI and SAFI combination that was not negotiated with the peer are now always silently discarded.
  - All ECMP paths for an IPv4 or IPv6 prefix must be labeled-unicast or unlabeled; a mix of path types is not supported. Similarly, the BGP FRR primary and backup paths for an IPv4 or IPv6 prefix must be the same type.
  - The **advertise-label** command, and the corresponding SNMP object, are obsoleted in Release 14.0.R4, which means that Nokia 5620 SAM and other SNMP management platforms must be appropriately updated to support the new CLI commands and SNMP objects. This requires an upgrade of the SAM software to the Nokia 5620 SAM 14.0.R5 release.

- The Nokia SR OS YANG modules have been reorganized to use submodules for the different areas of the SR OS configuration data model. There is only a single “module” (nokia-conf) which simplifies namespaces in requests and responses. This change is not backwards compatible with the Nokia YANG modules published in releases prior to Release 14.0.R4.
- If upgrading from any release to Releases 14.0.R4 and higher, Nokia highly recommends that the following configuration of the primary accounting policy be copied to the duplicate accounting policy to preserve the same accounting behavior.
  - accounting modes
  - update interval and jitter
  - session ID format
  - customer record
  - include attributes
  - account request script
  - tunnel format
- When upgrading to Release 14.0.R4 from an earlier release, there is a mandatory firmware upgrade for CPM5, CPM-X20, CPM-X16, CCM-X20 and CCM-e cards. During the software upgrade, the cards that require new firmware will automatically update their firmware when they are rebooted as part of the normal software upgrade process (ISSU or non-ISSU). The firmware update will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the cards are not removed while they are reprogramming the firmware.

  On 7950 XRS-40 systems, the upgrade to Release 14.0.R4 or higher can take up to 1 hour for all four CPMs/CCMs to reboot with their firmware upgrades. The CPMs on the Master chassis (CPMs A and B) will be in the “down” state during their firmware upgrade phases, while the CPMs on the Extension chassis (CPMs C and D) will be in the “provisioned” state during their firmware upgrade phases. [213165, 213169, 223695, 229675]
- Previously unsupported (but not blocked) OpenFlow GRT embedding should be removed from the IP/IPv6 filter associated with an R-VPLS interface. If OpenFlow embedding is not removed prior to ISSU, the filter association with the R-VPLS interface will be lost. [224266]
- Previously unsupported (but not blocked) action forward next-hop interfaces should be removed from the IP/IPv6 filter associated with an R-VPLS interface. If action forward next-hop are not removed prior to ISSU, the filter association with the R-VPLS interface will be lost. [229931]
- When upgrading to Release 14.0.R4 or higher, mixed-speed LAG with per-link-hashing enabled, newly introduced port mapping optimization may cause the links to be redistributed differently from previous releases. [236089]

9.1.10 Upgrading to Release 14.0.R3 or Higher

- Release 14.0.R3 changes the way XML Accounting files are formatted. Parsing functions in operator OSS layers may need to be adjusted if they had custom logic to work around the invalid SR OS XML formatting. Prior to Release 14.0.R3, the XML encoding used in SR OS accounting files for certain special characters was invalid. As of Release 14.0.R3, SR OS accounting files correctly encode the special characters as “&lt;”, “&gt;”, “&amp;”, “&apos;”, and “&quot;” instead of placing characters such as “<” directly into the accounting files. OSS parsing logic for Releases 14.0.R3 and higher XML Accounting files must be able to handle the standard XML encoding for the special characters.
- When upgrading to Release 14.0.R3 or higher from an earlier release, there is a mandatory firmware upgrade for CPM-e cards. During the software upgrade, the cards that require new firmware will automatically update their firmware when they are rebooted as part of the normal software upgrade process. The firmware update will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the cards are not removed while they are reprogramming the firmware. [229474]

9.1.11 Upgrading to Release 14.0.R1 or Higher

- The **tod-suite** and **time-range** commands should be removed from all configurations prior to upgrading to Release 14.0.R1 or higher. Management systems, such as Nokia NFM-P (formerly 5620 SAM), can provide similar functionality if required.

  For non-ISSU upgrades, and when executing scripts, the deprecated commands will be ignored and a message printed for each deprecated command, the command will be skipped and the execution will proceed the rest of the configuration. For upgrades, and when executing scripts, the deprecated commands will be ignored and a message printed for each deprecated command, the command will be skipped and the execution will proceed with the rest of the configuration.

  Ignoring the deprecated commands could, in specific circumstances, impact the execution of the remainder of the configuration. In certain configurations, it is possible to run out of resources after the deprecated commands have been ignored and the configuration has failed to load. The following is an example of such a circumstance:

  If there is a **tod-suite** provisioned with QoS policies applied but no time ranges, for example:

  ```
  tod-suite "tod" create
  egress
  qos 100
  ```

```
exit
ingress
qos 100
exit
exit
```

In this situation, any time the **tod-suite** is applied to a SAP, it will apply its ingress/egress QoS policies to that SAP permanently as there are no start/stop times. It renders whatever QoS policies have been applied directly on the SAP irrelevant, at least in terms of resources used. So with a SAP configured as follows:

```
sap 1/1/1:1 create
tod-suite "tod"
ingress
qos 200
exit
egress
qos 200
exit
exit
```

If SAP-ingress/egress QoS policies 200 would consume more resources than SAP-ingress/egress QoS policies 100, then, when booting up and ignoring all ToD configurations, it is possible that the resulting configuration would consume more resources than it would have with all of the **tod-suite** lines being executed. Hence, it is possible to run out of resources after ignoring the deprecated commands, and the configuration would fail to load.

9.1.12 Upgrading from Release 13.0.R5 to 13.0.R6 or Higher

- When upgrading from Release 13.0.R5 to Release 13.0.R6 or higher, there is a mandatory firmware upgrade for all CPMs and IOMs on the 7750 SR-a4/a8.

  During the software upgrade, the cards that require new firmware will automatically update their firmware when they are rebooted as part of the normal software upgrade process. The firmware update will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the cards are not removed while they are reprogramming the firmware. The Operational State of a card that is reprogramming its firmware will be displayed as “provisioned” under **show card** and the Equipped Type will be displayed as “not equipped”. [208437, 216782, 217615]

9.1.13 Upgrading to Release 13.0.R5 or Higher

- When upgrading to Release 13.0.R5 or higher from a previous release, there is a mandatory firmware upgrade for certain cards and platforms:
  - 7750 SR-a4/a8: all CPMs and IOMs
  - 7750 SR-7/12/12e: the CPM5 has new mandatory firmware in Release 13.0.R5

  During the software upgrade, the cards that require new firmware will automatically update their firmware when they are rebooted as part of the normal software upgrade process (ISSU or non-ISSU). The firmware update will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the cards are not removed while they are reprogramming the firmware. The Operational State of a card that is reprogramming its firmware will be displayed as “provisioned” and the Equipped Type will be displayed as “not equipped” in the output of the **show card** command.

9.1.14 Subscriber Management

- Due to increased memory requirements as a result of new software features, the maximum subscriber-host scale is 128k per system for the 7750 SR-7/12 and 7450 ESS-7/12 equipped with CPM3. This limit is not enforced by the system. For existing deployments that need a higher subscriber-host scale and want to upgrade to SR OS Release 13.0.R1 or higher, it is recommended to install CPM5 to provide more memory capacity. [199108]

9.1.15 Upgrading to Release 13.0.R1 or Higher

- Upgrading from a release earlier than Release 12.0.R2 to Release 13.0.R1 or higher can incorrectly change the configuration of **nat outside pool redundancy** to shutdown state. This configuration must be manually corrected. [215881]
- With the introduction of LDP IPv6 in Release 13.0.R1, a FEC for each of the IPv4 and IPv6 system interface addresses is advertised and resolved automatically by the LDP peers when the LDP session comes up, regardless of whether the session is IPv4 or IPv6.

  To avoid the automatic advertisement and resolution of IPv6 system FEC when the LDP session is IPv4, the following procedure must be followed before and after the upgrade to the SR OS version which introduces the support of LDP IPv6.

1. Before the upgrade, implement a global import prefix policy which rejects prefix [::0/0 longer] to prevent IPv6 FECs from being installed after the upgrade.
2. Standard software upgrade case:
   - If new IPv4 sessions are created on the node, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
   - On older, pre-existing IPv4 sessions, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
3. When *all* LDP IPv4 sessions have dynamic capabilities enabled, with per-peer FEC-capabilities for IPv6 FECs disabled, then the global import policy can be removed.

9.1.16 RMON

- RMON entries that referenced deprecated MIB entries are not automatically modified and re-saved with the MIB variable that may have replaced it. MIB variable changes are often due to a change in the indexing structure for such tables. Refer to the MIBs distributed with your SR OS image set and compare those as needed to MIBs from the prior SR OS release to identify changes and update the corresponding SNMP object or OID references in the configuration file.

9.1.17 MLD

- The checks for a valid link local address are corrected for some cases.

  Prior to Release 12.0.R4, addresses in the range of FE80::/10 were accepted (for example, FE81:: was accepted). In Releases 12.0.R4 and higher, the check is corrected and only addresses in the range of FE80::/64 are accepted.

  This will have an impact when performing an upgrade: configured values not in the FE80::/64 range will be rejected.

  Impacted configuration commands are:

  – **config**>**router**>**mld**>**group-interface** *group-interface-name*>**query-src-ip** *link-local address*

  – **config**>**service**>**vprn** *service-id*>**group-interface** *group-interface-name*>**query-src-ip** *link-local address*

  – **config**>**router**>**mld**>**grp-if-query-src-ip** *link-local address*

- **config**>**service**>**vprn** *service-id*>**mld grp-if-query-src-ip** *link-local address*
- **config**>**router**>**interface** *interface-name*>**ipv6 link-local-address** *link-local-address*
- **config**>**service**>**vprn** *service-id*>**interface** *interface-name*>**ipv6 link-local-address** *link-local-address* [172857]

9.1.18 Upgrading to Release 12.0.R1 or Higher

- A configuration with an IPv6 prefix present in the **router>router-advertisement interface** context on a non-mixed mode 7450 ESS will fail to execute from Release 12.0.R1 onward. It was possible in releases prior to Release 12.0.R1 to configure, although this was functionally not supported. If such a configuration exists, it has to be removed prior to upgrading to Releases 12.0.R1 and higher.
- The configuration command **configure system security user** *user-name* **console login-exec** " " (single space URL) will fail to execute from Release 12.0.R1 onward. Prior to Release 12.0.R1, it was possible to configure this, although it was not a valid URL. If such a configuration exists, it must be removed/updated prior to upgrading to Release 12.0.R1 or higher.
- The LFA SPF policy feature generalizes the use of admin-group and SRLG to non-MPLS interfaces. To that end, the definition of admin-groups and SRLGs has been moved from the **config>router>mpls** context to the new **config>router>if-attribute** context. The binding of MPLS interfaces to admin-group or SRLG remains under **config>router>mpls>interface**. When upgrading to Release 12.0.R1 or higher, all user-configured admin groups and SRLGs under the **config**>**router**>**mpls** context will automatically be moved under the new context.

9.1.19 DHCP

- When upgrading from Release 10.0.R10 through 10.0.R15 or from Release 11.0.R1 through 11.0.R7 to Release 12.0.R1 or higher, and DHCPv6 server and/or DHCPv6 relay on subscriber interfaces is/are enabled to assign IA_NA addresses, it may be required to add the global configuration parameter **adv-noaddrs-global esmrelay server** under the **config**>**system**>**dhcp6** context for backward compatibility. This parameter will send the “NoAddrsAvail” status code in DHCPv6 advertise messages at the global DHCP message level instead of at the default IA_NA option level.

9.1.20 Routing Policies

- In Releases 12.0.R1 and higher, the use of a community, **as-path**, **as-path-group** or **prefix-list** name starting and ending with “@” is no longer allowed. @...@ is used as identification for parameters being used in policies. Configuration files containing such names will fail to execute for upgrades from a release earlier than Release 12.0.R1 to Release 12.0.R1 or higher. [173346]

9.1.21 Upgrading to Release 11.0.R7 or Higher

- Starting with Release 11.0.R7, configuration changes are required for TACACS+ servers to authorize global commands. Global commands such as **info**, **exit**, and others, except the **logout** command, should be explicitly added to the configuration in the TACACS+ server. There are no changes required in the configuration on the SR OS node for this issue. A list of all global commands can be found in the *SR OS Basic System Configuration Guide*, or by entering **help globals** at the CLI prompt. [171214]

9.1.22 Upgrading from Release 11.0.R1 or 11.0.R2

- The parameter **port-forwarding-dyn-block-reservation** was introduced in Release 11.0.R1 and was incorrectly allowed to be configured for type L2-Aware NAT pools. In Releases 11.0.R3 and higher, a check was added to disallow the configuration of the parameter in combination with type L2-Aware NAT pools. Prior to upgrade, the parameter **port-forwarding-dyn-block-reservation** should be removed from the NAT configuration when having a type L2-Aware NAT-group configured. More details can be found in TA 13-1007. [163525]

9.1.23 LDP

- When upgrading from Release 11.0.R3, 11.0.R4, or 11.0.R5 to Releases 11.0.R6 and higher, the default setting for LDP event 2003 changed from generate to suppress. This value must be manually changed after the upgrade to properly save the newly corrected default setting of suppress. The default of suppress had been the default in Release 11.0.R2 and all prior releases. [170911]

9.1.24 Upgrading to Release 11.0.R4 or Higher on 7950 XRS-20

- The tmnxPortID mapping has changed for the 7950 XRS-20 platform. Refer to TIMETRA-TC-MIB for specific details.
- On upgrade, port indices in the SNMP MIB will not be preserved on these platforms. Management software that expects the old mapping may need to be updated.

9.1.25 R-VPLS

- Routed-VPLS (R-VPLS) does not support configuration of line card MAC filters. This restriction is now properly enforced starting with Releases 11.0.R1 and higher. A router using an SR OS version that enforces the restriction will not load a configuration that includes MAC filters in the context of R-VPLS. Before loading such a configuration either from a saved file or as part of an SR OS router upgrade, MAC filter configuration must be removed from the R-VPLS context.
- An R-VPLS service does not support Multicast-VLAN-Registration (MVR). This restriction is enforced in Releases 11.0.R1 and higher. With Release 10.0, it was possible to configure MVR options below a Routed-VPLS service. Before upgrading from Release 10.0, those options must be removed from the configuration, or loading the saved file will fail. [163006]

9.1.26 Filter Policy Consideration when Upgrading from Release 10.0.R4 or Higher to Release 11.0.R1 or Higher

- Starting with Release 11.0.R1, SR OS enforces the rule that a single CLI filter policy entry should not exceed the allowed hardware resources. Operators are advised to verify that a 10.0 configuration that uses match list in filter policies does not exceed the recommended limit prior to an upgrade. Failure to do so will result in configuration failure during an upgrade if the entry exceeds the enforced limits. The enforced rule allows 2048 hardware sub-entries per line card filter policy entry and 256 hardware sub-entries per CPM filter policy entry (approx. 25% margin atop Release 10.0.R4 recommended/supported limits. Refer to the 10.0 and 11.0 Release Notes, for Known Limitation 142472, for more information.

9.1.27 Upgrading to Release 11.0.R1 or Higher

- Support for the read-only radiusServerTable and read-only tacplusServerTable in the TIMETRA-SYSTEM-MIB has been removed in Releases 11.0.R1 and higher. The alternative readable and writable tables tmnxRadiusServerTable and tmnxTacPlusServerTable in the TIMETRA-SECURITY-MIB should be used instead. [131834]
- A new support.tim file has been introduced in Release 11.0.R1 as part of the SR OS software image package of *.tim files. All *.tim files should be copied together as a package when performing actions such as upgrades or backing up images. The support.tim file contains SR OS image data that is required for all platforms and configurations, and is not related to Nokia support services or the **admin tech-support** functionality.

  When upgrading from a release prior to Release 11.0.R1 to Releases 11.0.R1 and higher, the support.tim file must be manually synchronized (copied) across to the standby CPM. See step 5 of the Standard Software Upgrade Procedure. Releases prior to Release 11.0.R1 do not use the support.tim file, and hence the **synchronize** command will not copy it.

9.2 AA Signatures Upgrade Procedure

This section describes the AA Signatures Upgrade Procedure, which can be used to upgrade ISAs in 7750 SR-7/12/12e, 7750 SR-c4/c12 and ESS-7/12 to a new AA signature load without upgrading/impacting the router itself only when no firmware update is required.

If the above criterion does not apply, the Standard Software Upgrade Procedure must be performed.

Note:

- Although the software upgrade can be performed using a remote terminal session, Nokia recommends that the software upgrade procedure be performed at the system CONSOLE device where there is physical access to the 7750 SR or 7450 ESS as remote connectivity may not be possible in the event there is a problem with the software upgrade. Performing the upgrade at the CONSOLE with physical access is the best situation for troubleshooting any upgrade problems with the help of the Nokia Technical Assistance Center.
- This procedure applies to all ISA cards.

Step 1. Back up existing images and configuration files

New software loads may make modifications to the configuration file which are not compatible with older versions of the software.

Note:

- Configuration files may become incompatible with prior releases even if no new features are configured. The way in which a particular feature is represented in the configuration file may be updated by the latest version of the operating software. The updated configuration file would then be an unknown format to earlier software versions.

Nokia recommends making backup copies of the software image and configuration files (including bof.cfg and *.ndx persistency files). These backups will be useful in case reverting to the old version of the software is required.

Step 2. Copy Application Assurance ISA-AA.TIM file to cf3:

Application Assurance software and signatures are included in the isa-aa.tim file. This file must be copied to the same cf3: directory as the current SR OS images running on the router. It is good practice to place all of the image files for a given release in an appropriately named subdirectory off the root, for example, “cf3:\13.0.R1”.

As a result of this step, when upgrading the AA software only on an older SR OS software, the new isa-aa.tim file overwrites the existing software on the flash card.

Step 3. Synchronize boot environment

Active and standby CPM/CFM boot environments must be synchronized if the router has redundant CPM/CFMs.

– Use **admin redundancy synchronize boot-env** to synchronize the boot environments between the active and standby CPM/CFMs.

Step 4. Load new image for ISA card

After the boot environment has been synchronized, the new AA image needs to be loaded on the CPM/CFM.

– Use **admin application-assurance upgrade** to load the new isa-aa image on the CPM/CFM.

– Use **show application-assurance version** to verify new isa-aa image version running on the CPM/CFM.

– Use **show mda** to verify ISA card status.

```
A:ALU-ABC>show>app-assure# version
===========================================================================
Versions of isa-aa.tim in use
===========================================================================
CPM : TiMOS-M-13.0.R2
1/2 : TiMOS-M-13.0.R1
```

```
3/2 : TiMOS-M-13.0.R1
===========================================================================

A:router1# show mda
===========================================================================
MDA Summary
===========================================================================
Slot   Mda   Provisioned Type                  Admin Operational
               Equipped Type (if different)    State State
---------------------------------------------------------------------------
1      2     isa-aa                            up    ISSU/standby
               isa-ms
...
3      2     isa-aa                            up    ISSU/active
               isa-ms
===========================================================================
```

Step 5. Reset the ISA cards to load the new image

The ISA cards must now be reset to load the new image.

Note:

- The system does not allow cards to run in an ISSU state indefinitely; the system automatically resets the ISA cards after 2 hours. The “Comments” field in the **show card state** output displays the time until the system resets the ISA card in the ISSU state.

The timing and order of the ISA card resets should be sequenced to maximize the effectiveness of any redundancy. When redundancy is deployed, protecting (standby) ISA cards should be reset first, and admin activity switch should be forced first (**config card mda** *m*/*n* **shutdown**) before an active ISA card is reset.

- Use **shutdown mda** *m*/*n* to shut down an ISA card
- Use **clear mda** *m*/*n* to reset an ISA card
- Use **no shutdown mda** *m*/*n* to enable an ISA card
- Use **show application-assurance version** to verify the isa-aa signatures version loaded on the CPM/CFMs and the ISA cards

The sample output below shows the operational state transitions for a single Application Assurance group with one active and one protecting (standby) ISA card.

1. Before reset starts:

```
A:ALU-ABC>show>app-assure# version
===========================================================================
Versions of isa-aa.tim in use
===========================================================================
CPM : TiMOS-M-13.0.R2
1/2 : TiMOS-M-13.0.R1
```

```
3/2 : TiMOS-M-13.0.R1
===============================================================================

A:router1# show mda
===============================================================================
MDA Summary
===============================================================================
Slot   Mda    Provisioned Type                    Admin Operational
                Equipped Type (if different)      State State
-------------------------------------------------------------------------------
1      2      isa-aa                              up    ISSU/standby
                isa-ms
...
3      2      isa-aa                              up    ISSU/active
                isa-ms
===============================================================================
```

2. After the standby ISA card is reset and comes back up:

```
A:ALU-ABC>show>app-assure# version
===============================================================================
Versions of isa-aa.tim in use
===============================================================================
CPM : TiMOS-M-13.0.R2
1/2 : TiMOS-M-13.0.R2
3/2 : TiMOS-M-13.0.R1
===============================================================================

A:router1# show mda
===============================================================================
MDA Summary
===============================================================================
Slot   Mda    Provisioned Type                    Admin Operational
                Equipped Type (if different)      State State
-------------------------------------------------------------------------------
1      2      isa-aa                              up    up/standby
                isa-ms
...
3      2      isa-aa                              up    ISSU/active
                isa-ms
===============================================================================
```

3. After the ISA card activity switch (shutdown of active card to force activity switch):

```
A:ALU-ABC>show>app-assure# version
===============================================================================
Versions of isa-aa.tim in use
===============================================================================
CPM : TiMOS-M-13.0.R2
1/2 : TiMOS-M-13.0.R2
3/2 : TiMOS-M-13.0.R1
===============================================================================

A:router1# show mda
===============================================================================
MDA Summary
===============================================================================
Slot   Mda   Provisioned Type                  Admin Operational
               Equipped Type (if different)    State State
---------------------------------------------------------------------------
1      2     isa-aa                            up    up/active
               isa-ms
...
3      2     isa-aa                            down  ISSU/standby
               isa-ms
===========================================================================
```

4. After the newly inactive ISA card is reset, comes back up (**clear** command executed) and is re-enabled (**no shutdown** executed):

```
A:ALU-ABC>show>app-assure# version
===========================================================================
Versions of isa-aa.tim in use
===========================================================================
CPM : TiMOS-M-13.0.R2
1/2 : TiMOS-M-13.0.R2
3/2 : TiMOS-M-13.0.R2
===========================================================================

A:router1# show mda
===========================================================================
MDA Summary
===========================================================================
Slot   Mda   Provisioned Type                  Admin Operational
               Equipped Type (if different)    State State
---------------------------------------------------------------------------
1      2     isa-aa                            up    up/active
               isa-ms
...
3      2     isa-aa                            up    up/standby
               isa-ms
===========================================================================
```

Step 6. Update the AA policy and enable the new applications and protocol signatures

When the CPM/CFMs and ISA cards are using the latest image, update the AA policy definition and enable the new protocols available in this release. This process updates existing applications and corresponding app-filters maintained by Nokia, and creates newly supported applications.

- The operator must open a standard ticket, priority 3, to Nokia technical support, and provide a technical support file and the target AA software release deployed in the network.
- The technical support team will provide the following configuration update file to update the AA policy, to be executed on the target nodes: 7750# exec ftp://user:pass@ftp-server-ip/path/<aaconfig-delta-update-file-name>

9.3 Standard Software Upgrade Procedure

This section describes the Standard Software Upgrade Procedure that is service-affecting and must be used:

- when a manual firmware update is required (**admin reboot upgrade**).
- on routers with non-redundant CPM or CFM
- when ISSU is not supported in a given release

Each software release includes a BOOT Loader (boot.ldr). The BOOT Loader performs two functions:

1. Initiates the loading of the SR OS image based on the Boot Options File (bof.cfg) settings
2. Reprograms the boot ROM and firmware code on the CPM or CFM and IOM/IMM/ISM/XCM cards to the version appropriate for the SR OS image.

This section describes the process for upgrading the software and, if necessary, the boot ROM and firmware images with the BOOT Loader.

The software checks the firmware images on the CPM or CFM and IOM/IMM/ISM/XCM and reports any mismatch. If the loaded version is earlier than the expected version, the firmware may need to be upgraded; a console or log message will indicate if a firmware upgrade is required. If the firmware version loaded is later than the expected version, no firmware programming is required.

Note:

- Although the software upgrade can be performed using a remote terminal session, Nokia recommends that the software upgrade procedure be performed at the system CONSOLE device where there is physical access as remote connectivity may not be possible in the event there is a problem with the software upgrade. Performing the upgrade at the CONSOLE with physical access is the best situation for troubleshooting any upgrade problems with the help of the Nokia technical assistance center.
- Automatic firmware updates may occur for CPM and IOM/IMM/ISM/XCM cards running older firmware after an SR OS upgrade. The **clear card** command or physical removal of a card must not be performed until the card is operationally up after an SR OS upgrade. This procedure also applies when subsequently adding new IOMs/IMMs/ISMs/XCMs (that may have older firmware) to a chassis. An event log with “firmware upgraded” message will be issued if a firmware update had occurred for a card.

Step 1. Back up existing images and configuration files

New software loads may make modifications to the configuration file which are not compatible with older versions of the software.

Note:

- Configuration files may become incompatible with prior releases even if no new features are configured. The way in which a particular feature is represented in the configuration file may be updated by the latest version of the operating software. The updated configuration file would then be an unknown format to earlier software versions.

Nokia recommends performing an **admin save** and then making backup copies of the BOOT Loader (boot.ldr), software image and configuration files (including bof.cfg and *.ndx persistency files). These backups will be useful in case reverting to the old version of the software is required.

If Lawful Intercept (LI) is being used on the router and **bof li-local-save** is enabled, then the operator may want to save the LI configuration via **configure li save** and then backup the li.cfg file.

If the firmware version loaded is later than the expected version reported by the BOOT Loader, no firmware programming is required.

Step 2. Copy the SR OS images to cf3:

The SR OS image files must to be copied to the cf3: device on the active CPM or CFM (only on the master chassis for 7950 XRS-40). It is good practice to place all the image files for a given release in an appropriately named subdirectory off the root, for example, “cf3:\14.0.R3”. Copying the boot.ldr and other files in a given release to a separate subdirectory ensures that all files for the release are available should downgrading the software version be necessary.

Note:

- As of Release 11.0.R1, the support.tim file must also be copied for all platforms and configurations.

Step 3. Copy boot.ldr to the root directory on cf3:

The BOOT Loader file is named boot.ldr. This file must be copied to the root directory of the cf3: device of the active CPM/CFM (only on the master chassis for 7950 XRS-40).

Step 4. Modify the Boot Options File to boot the new image

The Boot Options File (bof.cfg) is read by the BOOT Loader and indicates primary, secondary and tertiary locations for the image file. The bof.cfg should be modified as appropriate to point to the image file for the release to be loaded. Use the **bof save** command to save the Boot Options File modifications.

Step 5. For Redundant CPMs or CFMs, synchronize boot environment

On systems with Redundant CPMs or CFMs, copy the image files and Boot Options File to the redundant CPM or CFM with **admin redundancy synchronize boot-env**.

When upgrading from a release prior to Release 11.0.R1 to Releases 11.0.R1 and higher, the support.tim file must be manually synchronized (copied) across to the standby CPM/CFM. Releases prior to Release 11.0.R1 do not use the support.tim file and hence the **synchronize** command will not copy it.

Step 6. Reboot the chassis

The chassis should be rebooted with the **admin reboot** command.

Step 7. Verify the software upgrade

Allow the boot sequence to complete and verify that all cards come online.

Software upgrade is successfully executed if the parsing of the configuration file completes as expected and there are no errors shown via a CONSOLE session or in the output of the **show boot-messages** CLI command.

If the configuration-file parsing stops with the error “CRITICAL: CLI #1002 The system configuration is missing or incomplete because an error occurred while processing the configuration file”, then check for known causes in the Release Notes or contact your Nokia support organization. Executing **admin save** at this point could result in the loss of the configuration.

To continue with the configuration-file parsing, remove the conflicting parameter from the loaded configuration file and re-execute it using the **execute** CLI command, or leave the loaded configuration file untouched and revert to the old version of the software.

Note:

- If any card fails to come online after the upgrade, contact the Nokia technical assistance center for information on corrective actions.

Nokia recommends saving the configuration with **admin save** after an upgrade has been performed and the system is operating as expected. This will ensure that all configuration is saved in a format that is fully compatible with the newly-running release.

10 Usage Notes

The following information supplements or clarifies information in the manuals for Release 15.1.R3 of SR OS.

Note:

- Usage notes added in this release are marked **[NEW]**.

10.1 Common Software Image Set for All Platforms

- A common software image set is used across the 7450 ESS, 7750 SR, and 7950 XRS platforms.

10.2 XCM and SFM Recovery Behavior

- In a 7950 XRS system, at least one SFM must be fully operational in order for the XCMs, XMAs and standby CPM to be in service. If there are no operating SFMs in the system, then the XCMs, XMAs and standby CPM will be held in a “booting” operational state.
- In a 7950 XRS system, at least one C-XMA/XMA in an XCM must be fully operational for the XCM to be in service. If there are no operating C-XMAs/XMAs in an XCM, then the XCM will be held in a “booting” operational state.

10.3 7750 SR-12e

- For optimal performance, Nokia recommends that up to four IOMs/IMMs for the 7750 SR-12e are installed in up to four consecutive slots (for example, slots 1-4 or 2-5).

10.4 7450 ESS-7/12 and 7750 SR-7/12/12e

- Specific engineering rules may apply when mixing FP2- and FP3-based line cards; contact your Nokia representative for further details.

10.5 Impedance Panels

- Impedance panels must be purchased and installed in all systems in which a line card is used. These impedance panels provide highly efficient air flow in support of the higher performing IOM3-XP/-B/-C, IOM4-e, IOM4-e-B and newer IMM/ISM modules. Even when only one IMM/IOM/ISM is deployed, impedance panels are required.

10.6 Multiservice Integrated Services Adapter (ISA)

The following tables list IOM and IMM support for ISA applications:

Table 23 **Compatible 7750 SR IOMs and IMMs for ISA Applications**

	IOM3-XP/-b/-c	**MS-ISM/MS-ISA2 IMM**	**MS-ISM-E MS-ISA2-E IMM**	**MS-ISA2 on IOM4-e and IOM4-e-B**	**MS-ISA2-E on IOM4-e and IOM4-e-B**	**MS-ISA2 on IOM-e (SR-1e/2e/3e)**	**MS-ISA2-E on IOM-e (SR-1e/2e/3e)**
Application Assurance (isa-aa/isa2-aa) [1], [2]	✓	✓	✓	✓	✓	✓	✓
Retransmission, Fast Channel Change, and Video Quality Monitoring (isa-video/isa2-video)	✓	✓	✓	✓			
Tunnel Services, including IPsec (isa-tunnel/isa2-tunnel) [1]	✓ [3]	✓		✓		✓	

Table 23 **Compatible 7750 SR IOMs and IMMs for ISA Applications**

	IOM3-XP/-b/-c	MS-ISM/MS-ISA2 IMM	MS-ISM-E MS-ISA2-E IMM	MS-ISA2 on IOM4-e and IOM4-e-B	MS-ISA2-E on IOM4-e and IOM4-e-B	MS-ISA2 on IOM-e (SR-1e/2e/3e)	MS-ISA2-E on IOM-e (SR-1e/2e/3e)
Network Address Translation (isa-bb/isa2-bb) [1]	✓	✓	✓	✓	✓	✓	✓
L2TP LNS Service (isa-bb/isa2-bb)	✓	✓	✓	✓	✓	✓	✓
WLAN-GW (isa-bb/isa2-bb)	✓	✓ [4]	✓ [5]	✓ [6]	✓ [6]		

Notes:

1. Application Assurance, Tunnel and IPsec services and NAT are also supported on the 7750 SR-c12.
2. Application Assurance is also supported on the 7750 SR-c4.
3. MS-ISA only. Not supported on MS-ISA-E.
4. MS-ISM only. Not supported on IMM.with a single MS-ISA2.
5. MS-ISM-E only. Not supported on IMM with a single MS-ISA2-E.
6. Requires both MDA slots in the IOM4-e to be equipped with MS-ISA2 (-E) cards.

Table 24 **Compatible 7450 ESS IOMs and IMMs for ISA Applications, without Mixed Mode**

	IOM3-XP/-b/-c	MS-ISM/ MS-ISA2 IMM	MS-ISM-E/ MS-ISA2-E IMM	MS-ISA2 on IOM4-e and IOM4-e-B	MS-ISA2-E on IOM4-e and IOM4-e-B
Application Assurance (isa-aa/isa2-aa)	✓	✓	✓	✓	✓

Table 25 **Compatible 7450 ESS IOMs and IMMs for ISA Applications, with Mixed Mode**

	IOM3-XP/-b/-c	MS-ISM/ MS-ISA2 IMM	MS-ISM-E/ MS-ISA2-E IMM	MS-ISA2 on IOM4-e and IOM4-e-B	MS-ISA2-E on IOM4-e and IOM4-e-B
Application Assurance (isa-aa/isa2-aa)	✓	✓	✓	✓	✓
Tunnel Services, including IPsec (isa-tunnel/isa2-tunnel)	✓[1]	✓		✓	
Network Address Translation (isa-bb/isa2-bb)	✓	✓	✓	✓	✓
L2TP LNS Service (isa-bb/isa2-bb)	✓	✓	✓	✓	✓

Note:

[1] MS-ISA/ISA2 only. Not supported on MS-ISA-E/ISA-E.

10.7 Compact Flash Devices

- Only Nokia-sourced compact flash devices for the SR OS are supported.
- In Releases 13.0.R1 and higher, Nokia recommends that the compact flash in the CF3 slot be at least 2 Gbytes. The extra compact flash space is intended to support customers who may want to keep more than one copy of the software.
- Nokia recommends using cf1: or cf2: for event logs and dynamic data persistency.

10.8 Hardware

- SFPs with bad checksums cause traps and log events. The port will be kept operationally down with SFPs that fail to read or have invalid checksums. [62458]
- When a dual-rate SFP is connected to a GigE LX SFP, the auto-negotiation parameter must be turned off in order to get a link. [67690]
- The SR OS routers support qualified pluggable optic modules only. Refer to the current Nokia price list for supported modules. Third-party optics are not supported.

10.9 System

- When creating a new log file on a compact flash disk card, the system will check the amount of free disk space and the amount must be greater than or equal to the lesser of 5.2 MB or 10% of the compact flash disk capacity.
- SNMPv3 user authentication and privacy keys in the **config system security user** *user-name* **snmp authentication** command must be entered as maximum length strings. [18314]
- Manual editing of SNMP persistent index files can cause errors in loading the configuration file. Persistent index files should only be created by the system. [24327]
- When nodes are run in FIPS-140-2 mode (where only FIPS-140-2 algorithms are enabled and allowed), Nokia recommends only enabling the FIPS-140-2 mode on newly deployed nodes. Changing to FIPS-140-2 mode on live nodes should be avoided as there may be conflicts with existing configurations that are not consistent when running the node in FIPS-140-2 mode. Before enabling a pre-configured node to run in FIPS-140-2 mode, ensure all configurations in the configuration file are devoid of conflicting configurations that are not allowed in FIPS mode, such as the use of any unapproved cryptographic algorithms or certificates that are signed with unapproved algorithms. Refer to the *Basic System Configuration Guide* for details.
- If log 99 on the active CPM shows “Class CPM Module: failed, reason: Inactive CPM BOF LI config invalid” in a High-Availability setup, it indicates that:
  - the **bof li-separate** command has been issued
  - the standby CPM had experienced a reset (for example, **admin reboot standby**) and is currently in operation down state

  To restore the standby CPM and for **li-separate** to take effect, perform the following:

1. Ensure the bof.cfg file matches between the active and standby CPMs. If the bof.cfg does not match, update the bof.cfg on the active CPM, as the standby CPM is operationally down.
2. Back up both the bof.cfg and the configuration.
3. Issue the **admin reboot** command to reboot the chassis.

Note: The **li-separate** command always require a mandatory reboot of the chassis.

10.10 Satellites

- LLDP will now be automatically configured on host ports bound to satellite uplinks and will no longer configurable on host ports. Host ports with LLDP enabled cannot be bound to a satellite. Deprecated configuration executed via an older configuration file will be skipped with a message provided.

10.11 Multi-Chassis Synchronization

- Nokia recommends using the same CPM types in both chassis of the redundant MCS pair for production deployments. Different CPM types can be used during a hardware upgrade procedure.

10.12 NETCONF/YANG

- The following YANG modules are published and distributed as part of an SR OS image in the cflash/support directory:
  - Base-R13 SR OS YANG modules for configuration
  - Nokia YANG modules for state

  The Nokia YANG modules for configuration are available upon request. Please contact your Nokia representative.

10.13 Telemetry/gRPC

- The SR OS gRPC Telemetry interface in Release 15.1.R1 is based on OpenConfig gnmi.proto version 0.4.0. The gNMI specification is continuing to evolve leading up to a '1.0' version and future releases of SR OS are expected to implement later versions of gnmi.proto. Before upgrading to SR OS software releases that contain updated versions of gnmi.proto, clients/collectors must be updated to account for telemetry interface changes.

10.14 ATM

- 7750 SR and 7450 ESS in mixed mode allow configuration of user traffic on reserved ATM Forum UNI specification VCI values (VCIs from 0 to 31 inclusive). Nokia recommends not configuring any user traffic on those VCIs on any VP as other equipment may treat that traffic per the defined usage reserved to a given VCI value. Additionally, users must not configure VCIs 0, 3, 4, 6, and 7 on any VPI for services on ASAP MDAs, as those VCIs are exclusively used for their ATM Forum defined and reserved functionality. [53205]

10.15 MLPPP

- When a MLPPP bundle is out of service (oos), the Oper MTU and Oper MRRU are derived from the configured MRRU.
- Currently, LCP echo ids from 0–255 are separated into two ranges:
  - 0–127 is used for keepalive function
  - 128–255 is used for differential delay detection.

  Keepalive statistics only count echo packets with IDs from 0-127.
- In order to interoperate with other vendors’ MLPPP implementations, the MLPPP sub-layer will accept packets with or without leading zeros in the protocol field even though the 7750 SR and 7450 ESS in mixed mode do not advertise the protocol field compression (PFC) option during LCP negotiation. [25996, 29923]

10.16 APS

- Nokia recommends that the **lb2er-sd** and **lb2er-sf** alarms be enabled for SONET/SDH ports belonging to APS groups to better understand some APS group switchovers between the working and protect circuits.
- For SONET/SDH ports belonging to APS groups that have a very large difference in the transmission delay between the working and protect circuits, Nokia recommends that the hold down timers be increased from their default values.
- Increased APS group scaling (above 32 MC-APS and 64 SC-APS) requires CPM3 or higher for optimal switchover performance during failures affecting multiple groups. Nokia recommends CPM3 or higher for APS group scaling over 64 groups.

10.17 TCP Authentication Extension

- Keychains with no active entries will keep LDP and BGP peerings down. [57917]

10.18 Routing

- Nokia recommends that the preference value for BGP routes be set to a higher value than that of the internal (IGP) routes used to resolve the next-hop addresses of IBGP routes or routing instability can occur while the BGP routes are constantly re-learned. [31146]
- Any changes to multi-stream S-PMSI policy or a more preferred multi-stream S-PMSI (less or equal to current policy index) might cause a traffic outage; as such, it is recommended for any changes to multi-stream S-PMSI policies to be performed in a maintenance window.

10.19 Disallowed IP Prefixes

- The following IP address prefixes are not allowed by the unicast routing protocols and the Route Table Manager and will not be populated within the forwarding table:
  - 0.0.0.0/8 or longer

  - 127.0.0.0/8 or longer
  - 224.0.0.0/4 or longer (used for multicast only)
  - 240.0.0.0/4 or longer

  Any other prefixes that need to be filtered can be filtered explicitly using route policies.

10.20 IS-IS

- The granularity of the IS-IS hold timer is accurate only to within +/- 0.5s, so having a computed holdtime value of less than 2s may result in adjacencies being randomly dropped. Nokia recommends that **hello-interval**s and **hello-multiplier** values be adjusted accordingly, paying specific attention to the smaller hold-times computed on DIS systems. [29490]
- IS-IS authentication is not activated at any given level or interface unless both the authentication key and type are added at that level. For instance, if **hello-authentication-type** is set to password for an interface, it is not activated until a key is added at the interface level. [34256]

10.21 IS-IS TE

- The protocol sends advertisements with the IS-IS Traffic Engineering (TE) Router ID TLV when traffic engineering is disabled. [17683]

10.22 Auto-derived Route-Distinguisher (RD) in services with multiple BGP families

- In a VPLS service, multiple BGP families and protocols can be enabled at the same time. When **bgp-evpn** is enabled, **bgp-ad** and **bgp-mh** are also supported. It is important to note that a single RD is used per BGP instance and not per BGP family/protocol. The following rules apply.
  - The VPLS RD is selected based on the following precedence:
    - manual-RD or auto-RD always take precedence when configured
    - if there is no manual-RD/**auto-rd** configuration, the RD is derived from the **bgp-ad>vpls-id**

        - if there is no manual-RD/**auto-rd**/**vpls-id** configuration, the RD is derived from the **bgp-evpn>evi**, except for **bgp-mh**, which does not support evi-derived RD.
        - if there is no manual-RD**auto-rd**/**vpls-id**/**evi** configuration, there is no RD, and thus the service will fail
    - The selected RD (see above rules) will be shown in the “Oper Route Dist” field of the **show service id** *service-id* **bgp** command.
    - The service supports RD changes dynamically; for instance, the CLI allows the vpls-id to be changed even while it is being used to auto-derive the service RD for **bgp-ad**, **bgp-vpls** or **bgp-mh**. Note that, when the RD changes, the active routes for that VPLS will be withdrawn and re-advertised with the new RD.
    - If one of the mechanisms to derive the RD for a given service is removed from the configuration, the system will select a new RD based on the above rules. For example, if the **vpls-id** is removed from the configuration, the routes will be withdrawn, the new RD selected from the **evi**, and the routes re-advertised with the new RD.
    - Because the **vpls-id** takes precedence over the **evi** when deriving the RD automatically, adding **evpn** to an existing **bgp-ad** service will not impact the existing RD—this is important to support **bgp-ad** to **evpn** migration.

10.23 BGP

- Nokia recommends that the local address be configured when a router has multiple BGP peers to the same node. [113614]
- The static black-hole route should be created prior to receiving routes or creating the policy in combination with auto-bind-tunnel GRE. [160617]
- The VPN route selection for VPN routes that are not imported into any VPRN has been changed to take the IGP cost into account in Release 15.0.R1 and later, with the introduction of BGP optimal route reflection (ORR). In prior releases, the IGP cost was not considered for such routes to be selected as best route. [262800]

10.24 BGP Auto-Discovery

- On the 7450 ESS without mixed mode, only the L2-VPN address family is supported by BGP. This address family is used for BGP Auto-discovery for VPLS. Any commands or options for other address families in BGP or in routing policies are not supported on the 7450 ESS except in mixed mode.

10.25 BGP VPWS

- When a provisioned SDP that is used for a spoke-SDP is shut down, or there is a local LSP failure (causing the spoke-SDP to go down), a BGP-VPWS update will be sent to the adjacent PE with the CSV bit set to one. This, however, does not cause the spoke-SDP, site or SAP to go down on the adjacent PE. If the adjacent PE is the designated forwarder of a pair of dual-homed PEs, no designated forwarder failover occurs. The above situation can result in the designated forwarder being one of the dual-homed PEs but the remote PE using its pseudowire to the other dual-homed PE.

10.26 MPLS/RSVP

- The current bypass binding selection logic for Releases 7.0 and higher is the following:
  - For non-strict environment
    - a) Manual CSPF disjoint bypass
    - b) Manual CSPF !disjoint bypass
    - c) Dynamic CSPF disjoint bypass
    - d) Dynamic CSPF !disjoint bypass
  - For strict environment
    - a) Manual CSPF disjoint bypass
    - b) Dynamic CSPF disjoint bypass

  The above binding order has two collateral/detrimental effects when the non-strict option is selected:

  1. In presence of a disjoint Dynamic Bypass, a non-disjoint Manual Bypass may be selected instead.
  2. Non-CSPF Manual Bypass will never be selected. [66005]

- The enabling or disabling of Diff-Serv on the system requires that the RSVP and MPLS protocols be shut down. When first created in Release 7.0 or higher, RSVP and MPLS will be administratively down. The user must execute the **no shutdown** command for each protocol once all parameters under both protocols are defined. When saved in the configuration file, the **no shutdown** command is automatically inserted under both protocols to ensure they come up after a node reboot. In addition, the saved configuration file is organized so that all LSP-level and LSP path-level configuration parameters are executed after all MPLS and RSVP global- and interface-level parameters are executed.
- LSP MTU negotiation for P2MP LSP is not supported. End-to-end MTU along the S2L path needs to be large enough to support data traffic. [74835]

10.27 LDP

- On LDP interfaces and **targeted-session keepalive** commands, Nokia recommends that the **factor** setting be set to a value greater than 1 or it may lead to unexpected drops in LDP peerings. [67153]
- When a per-peer import/export policy, which is either non-existing, incorrectly configured or not committed yet is configured, it may result in the system rejecting any FEC from being imported/exported. The workaround is to ensure that the configuration files do not contain policy mis-configurations or mismatches between LDP and the policy manager.

10.28 IP Multicast

- If an **rp static-address** is configured, the current PIM implementation will install an implicit deny-all for 224.0.0.0/4. To re-permit this address range, another static entry for this range must be installed. [38630]
- MoFRR for PIM interfaces should be enabled on a hop-by-hop basis to ensure optimal MoFRR recovery.
- If auto-rebalancing is enabled, re-balancing when a new path becomes available is performed for active joins.
- Optimized IP-multicast replication over RSVP-TE spoke-SDPs using configurable multicast network domains requires all spoke interfaces to be configured exclusively on physical ports, LAG ports, or APS-protected ports. If that is not the case, the default replication will take place.

- To execute **mtrace** and **mstat** with protocol-protection enabled (**config>security>cpu-protection**), IGMP must be enabled on incoming interfaces. [160402]

10.29 PIM

- To ensure proper GRT/VRF extranet functionality, it is strongly recommend to shut down PIM inside the VPRN (**config>service>vprn>pim>shutdown**) when enabling **grt-extranet** functionality in this VPRN under the following cases:
  - enabling **grt-extranet** for the first time in the VPRN
  - configuring **grt-extranet group-prefix any** or **grt-extranet group-prefix 224.0.0.0/4**
  - configuring **grt-extranet group-prefix** for a group that is already present in the VPRN.

  To ensure proper per-group map extranet functionality, it is strongly recommend to shut down PIM inside the receiver VPRN (**config**>**service**>**vprn**>**pim**>**shutdown**) when enabling the per-group mapping extranet functionality in this VPRN under the following cases:
  - enabling per-group mapping for the first time in the VPRN (that is, configuring the first map entry)
  - configuring **group-prefix 224.0.0.0/4** inside the map (that is, mapping all multicast groups to one core instance). [186280]

10.30 QoS

- By default, the CBS value of newly-created queues in queue-group policies is zero percent. Adding queue-groups or other configuration may result in reservation of all available buffer space (CBS) so that there is no shared buffer space available and queues with CBS of zero percent will drop traffic. Expedited traffic for newly-created queues in queue-group policies with default CBS of zero percent may also be lost when there is congestion of non-expedited traffic. To prevent the loss of traffic, Nokia recommends that the CBS value be changed to at least one percent for expedited and non-expedited queues, or for non-expedited queues, to ensure that shared buffer space is available. Buffer memory can be monitored with the **show pools** command. [86843]

- On the 7750 SR-a4/a8, ingress multipoint traffic is forwarded using shared queuing instead of the multipoint shared queuing. Specifically, the first pass through the FP uses the regular service queues and the second pass uses the default shared unicast queues instead of the default shared multipoint queues. Consequently, any parameter changes (for example, rates and MBS/CBS) applied to the default shared multipoint queues will not have any effect on the received multipoint traffic. [184678]
- **profile-mode** queues in FP3 platforms use two offered statistic counters as opposed to four in non-FP3 platforms. This means FP3 unicast **profile-mode** queues provide offered-uncolored and a combined in-/out- profile offered-colored statistics. FP3 multicast **profile-mode** queues provide a combined offered-combined statistics and an offered-mcast-managed statistics for managed multicast. Starting in Release 10.0.R1, multicast **profile-mode** queues on non-FP3 platforms report offered-uncolored and offered-managed using separate counters. No new MIB object is added as part of these statistics changes. Since existing MIB objects are used, non-FP3 **profile-mode** multicast queue offered-managed and offered-uncolored are accounted using the same MIB object. The **show** command output displays offered-managed and offered-uncolored as separate statistics for **profile-mode** non-FP3 multicast queues. The **show** command output also displays different statistic counters based on platform type.

10.31 Filter Policies

- Starting with Release 11.0.R1, the maximum number of filter policies and filter policy entries per system is larger than the line card limit. Since filter statistics are maintained on line cards and aggregated on the CPM/CFM, when an entry is deleted from a given line card (that is, an entry is deleted, or a given filter policy is no longer used on a given line card), the CPM/CFM resets that entry’s counters to zero. If the counters are required, they should be retrieved prior to such a configuration change.
- Nokia recommends against deploying the same filter policy on both ingress and egress because ingress and egress filter policies support different functionalities (actions and/or match criteria).
- Using a filter policy on a line card or in a direction that does not support a given match criterion may result in an unexpected match by the filter entry. It is recommended to avoid such configurations.
- When a filter policy is used on a line card that does not support a given action or in a direction that does not support that action, the action is ignored; if the packet matches the entry, default action is executed.

- Starting from Release 11.0.R1, all newly-introduced filter policy functionality is no longer supported in combination with ToD functionality. Nokia recommends against configuring a filter policy that has both ToD and Release 11.0.R1 or newer filter policy enabled.

10.32 Services General

- Starting in Release 10.0.R3, a PW port needs to be created first (with **encap-type dot1q** or **qinq**) before it can be bound to the SDP. Configurations containing PW-port entries from releases prior to Release 10.0.R3 are not compatible. [134086]
- In Releases 15.0.R4 and higher, these objects have an optional *name* parameter on the **create** line:
  - all services (**configure service vprn**, **vpls**, **epipe**, etc)
  - **mirror-dest**
  - **configure service pw-template** contexts
  - **configure service customer**
  - **configure filter ip-filter**, **ipv6-filter** and **mac-filter**
  - **configure qos network**

  For example:
  - **configure service vprn** *service-id* [**name** *name*] **customer** *x* **create**
  - **configure filter ip-filter** *filter-id* [**name** *name*]
  - **configure qos sap-ingress** *policy-id* [**name** *name*]
- Although the CLI allows the user to configure any value, the **source-bmac** address being used in a B-VPLS service must not overlap with any configured static-MAC address or OAM MAC address in the same B-VPLS service.
- In Releases 15.1.R1 and higher, the **create** line *name* parameter (introduced in Release 15.0.R4) for the following objects can no longer be changed after the object is created (*name* is immutable):
  - all services (**configure service vprn**, **vpls**, **epipe**, etc)
  - **mirror-dest**
  - **configure service pw-template** contexts
  - **configure service customer**
  - **configure filter ip-filter**, **ipv6-filter** and **mac-filter**
  - **configure qos network**, **sap-ingress**, **sap-egress**

  For example:

- **configure service vprn** *service-id* [**name** *name*] **customer** *x* **create**
- **configure filter ip-filter** *filter-id* [**name** *name*]
- **configure qos sap-ingress** *policy-id* [**name** *name*]

In addition, the **service-name**, **pw-template-name**, **customer-name**, **filter-name**, and **policy-name** commands (inside the various **service**, **filter** and **qos** policies) no longer exist. They have been replaced by the *name* parameter on the **create** line.

See Changed or Deprecated Commands in Release 15.1.R1-1 for more information.

10.33 Proxy-ARP/ND recommended settings

When enabling Proxy-ARP/ND in a VPLS service, Nokia recommends the following configuration for the correct network behavior:

- Nokia recommends enabling **dynamic-arp-populate** or **dynamic-nd-populate** only in networks with a consistent configuration of this command in all PEs. In EVPN networks where some nodes do not support this feature, **dynamic-arp-populate** and **dynamic-nd-populate** should only be enabled if the EVPN nodes always advertise IP->MAC pairs in MAC routes. For example, when an SR OS router is used as a Data Center (DC) Gateway for a Nuage DC, the user should enable **dynamic-arp-populate** only if all the Nuage Vports in the service are type host or VM (since their IPs will be advertised in MAC routes).
- When using **dynamic-arp-populate/dynamic-nd-populate**, the **age-time** value should be configured to a value equal to three times the **send-refresh** value. This will help reduce the EVPN withdrawals and re-advertisements in the network.
- In case of large **age-time** values, it would be sufficient to configure the **send-refresh** value to half of the Proxy-ARP/ND age-time or FDB age-time.
- In scaled environments (with thousands of services) it is not recommended to set the **send-refresh** value to less than 300 seconds. In such scenarios, Nokia recommends using a minimum Proxy-ARP/ND **age-time** and FDB age of 900 seconds.
- The use of the following commands reduces or suppresses the ARP/ND flooding in an EVPN network, since EVPN MAC routes replace the function of the regular data plane ARP/ND messages:
  - **no garp-flood-evpn**
  - **no unknown-arp-request-flood-evpn**
  - **no unknown-ns-flood-evpn**

  - **no host-unsolicited-na-flood-evpn**
  - **no router-unsolicited-na-flood-evpn**

  Nokia recommends using these commands only in EVPN networks where the CEs are routers directly connected to an SR OS node acting as the PE. Networks using aggregation switches between the host/routers and the PEs should flood GARP/ND messages in EVPN to make sure the remote caches are updated and BGP does not miss the advertisement of these entries.

- When the **anti-spoof-mac** is used with Proxy-ARP/ND, ingress filters (in the access SAPs/SDP-bindings) should be configured to drop all traffic with destination **anti-spoof-mac**. The same MAC should be configured in all PEs where *dup-detect* is active.
- When Proxy-ND is used, the configuration of the following commands should be consistent in all the PEs in the network:
  - **router-unsolicited-na-flood-evpn**
  - **host-unsolicited-na-flood-evpn**
  - **evpn-nd-advertise**

  Since EVPN does not propagate the “router” flag in IPv6->MAC advertisements, in a mixed network with hosts and routers, if **evpn-nd-advertise** router is configured, unsolicited host NA messages should be flooded so that the entire network gets to learn all of the host and router ND entries. In the same way, **evpn-nd-advertise** host should be configured if unsolicited router NA messages are flooded.

10.34 Subscriber Management

- Dynamic data persistency (subscriber management, DHCP server, Python-policy cache, NAT port forwarding, Application Assurance or ANCP) usage notes are as follows.
  - Nokia recommends discontinuing the use of 256M and 1G compact flash cards for dynamic data persistency applications; using a 4G or 8G compact flash card is recommended. In Releases 13.0.R1 and higher, Nokia recommends using an 8G compact flash card when enabling multiple dynamic data persistency applications.
  - Dynamic data persistency should not be configured to use compact flash cards formatted with the Reliance file system.
  - Nokia recommends a maximum of two applications on the same compact flash card when using multiple dynamic data persistency applications.
  - CF3 must not be used as the location for dynamic data persistency.

  - XML accounting (stored on compact flash) should not be used in conjunction with dynamic data persistency. Nokia recommends RADIUS accounting as an alternative. [50940]

- Starting with Release 11.0.R1, a RADIUS server configured under the routing instance (base, management or VPRN service) **radius-server** context can be used for authentication and accounting applications simultaneously. It is now possible to configure an **auth-port** and an **acct-port** for each server. When upgrading from a release prior to Release 11.0.R1, the single port configured for the server is automatically migrated to the new configuration. In this case, both **auth-port** and **acct-port** will have the same value. This is not a problem for the active configuration, but needs to be manually updated if the server is used for multiple applications.
- A PPPoE session will no longer be automatically terminated by the system in the following cases:
  - Starting with Release 14.0.R1, the system will no longer terminate a local user database (LUDB)-authenticated PPPoE session when the LUDB configuration changes during the lifetime of the session.
  - Starting with Release 14.0.R2, the system will no longer terminate a PPPoE session when the DNSv4/NetBios name server information is updated via a local DHCP client renew.

  To update the PPPoE session in these cases it can be restarted via CLI or AAA instead.
- DHCPv6 server DUID configuration guidelines in multi-chassis redundancy scenarios are as follows:
  - In a redundant DHCPv6 server configuration, each server must have a unique DUID (configured as **server-id** in the **router** and **service vprn dhcp6 local-dhcp-server** CLI context). Configuring an identical DUID with failover mode **local** or **remote** can result in unpredictable or multiple prefix allocation.
  - In a multi-chassis redundant DHCPv6 proxy-server configuration, both proxy-servers must share the same DUID (configured as **server-id** in the **group-interface ipv6 dhcp6 proxy-server** CLI context). Configuring a different DUID can result in ignoring the lease renewal and release after an SRRP switchover.
- Configured values for **valid-lifetime** must be greater than **preferred-lifetime**. The CLI context does not check this. If configured **valid-lifetime** is less than **preferred-lifetime**, default values are used. [250467]

10.35 Use of BGP-EVPN, BGP-AD and BGP-MH in the same VPLS service

- BGP-EVPN, BGP-AD and BGP-MH (one site) can all be configured in the same VPLS service. If that is the case, the following considerations apply:
  - The configured BGP route-distinguisher and route-target are used by BGP for the two families (that is, EVPN and L2-VPN). If different import/export route targets are used per family, vsi-import/export policies must be used.
  - The **pw-template-binding** command under BGP does not have any affect on EVPN or BGP-MH. It is only used for the instantiation of the BGP-AD spoke-SDPs.
  - If the same import/export route-targets are used in two redundant systems for BGP-EVPN and BGP-AD, a VXLAN binding, as well as a FEC129 spoke-SDP binding, may be attempted between the two systems, creating a loop. If that is the case, the SR OS will allow the establishment of an EVPN VXLAN binding and an SDP-binding to the same far-end, but it will keep the SDP-binding operationally down. Only the VXLAN binding will be operationally up. [170951]

10.36 VPRN/2547

- A route policy statement entry referencing a non-existent prefix list, community list, or AS path list will be accepted without a warning when committing a route policy configuration. This kind of missing reference can be seen when executing **show router policy-edits**. [60879, 84264, 86129]

10.37 VXLAN

- VXLAN IPv6 packets are always transmitted with a zero UDP checksum as recommended by RFC7348 (VXLAN). This may cause issues in deployments where VXLAN IPv6 packets are encapsulated in IPsec. The packets will be dropped if the IPsec Gateway checks the UDP checksum on the private interface before adding the IPsec encapsulation, as required by RFC 2460. Note that RFC 6935 relaxes this requirement and allows IPsec Gateways to transmit zero UDP checksum packets received on their private interfaces. [264804]

10.38 IPsec

- IKE traffic should be treated as higher priority than any data plane traffic (like ESP) on the end-to-end path from a remote IPsec peer to a 7750 SR, which means that appropriate ingress/egress QoS policy should be configured on the corresponding network facing port (or SAP) and public tunnel-SAP of 7750 SR and any other network forwarding node along the way.
- CRL NUMBER is a non-critical CRL extension; the CRL file provisioned in **ca-profile** should not mark this extension as critical.
- Certificate configured in **cert-profile** should be an end-entity certificate; a CA certificate should not be configured in these places.

10.39 IPsec Compatibility

- The following tables list software and hardware tested for compatibility with IPsec services:

Table 26 **Compatible Devices for Dynamic LAN-to-LAN IPsec Tunnels**

Device	Tested Version
Nokia VPN Firewall Brick 1200	9.1
Bintec Funkwerk R1200WU	7.5 Rev 3

Table 27 **Compatible IPsec Soft Client**

Soft Client	Tested Version(s)
Cisco VPN Client	5.0.03.0560
Racoon	NetBSD running ipsec-tools 0.7
SafeNet SoftRemote	10.8.3
Shrewsoft	2.1.2
Strongswan	2.8.x, 4.2.x, 5.0.1

10.40 Mirror Service

- CLI commands entered under the **debug mirror-source** sub-menu are now automatically synchronized with the standby CPM/CFM. These commands must no longer be placed in the CLI script file that is executed with the **switchover-exec** command. [105122]

10.41 OpenFlow

- H-OFS supports statistics collection per entry for Flow Table and Logical Port Table. Due to large H-OFS scale, Nokia recommends that a single statistics request message from the controller does not map (using a wildcard or cookie) to more than 1000 Flow Table entries per cookie context per message or 10 Logical Port Table entries per message.

10.42 Application Assurance

- Operators using applications maintained by Nokia for analytics, charging, or control should update both protocol signatures and the AA policy definition on a regular basis. New and updated protocols are available in the isa-aa.tim file while the AA policy update is provided through Nokia technical support. See AA Signatures Upgrade Procedure for more details.
- The isa-aa.tim image is available in the same directory as other .tim images. The image contains the Application Assurance software used on MS-ISA and the protocol list loaded by the CPM. The Application Assurance software can be upgraded independently of the SR OS software within a major release of the SR OS.
- When an Application-Assurance group **dual-bucket-bandwidth** policer is configured, the default configuration will cause all packets to be dropped. Ensure that the **dual-bucket-bandwidth** policer is configured appropriately. [86311]
- Only properly negotiated TCP sessions are eligible for TCP performance sampling.
- Changes to the TCP performance sampling rates will only affect new traffic flows.

- The bandwidth capacity for an AA-subscriber is equal to the full capacity of the MS-ISA or MS-ISA2 card, provided there is a realistic diversity of traffic sessions. The bandwidth capacity of an individual traffic session is limited by the in-order analysis and the amount of high-touch processing required by each packet in the session.
- If a Forwarding Path (FP) is configured with one MDA type of ISA-AA and any other MDA type (except a second ISA-AA) on an IOM3 or on a 7750 SR-c4/c12 system, then the FP buffer allocation must be modified from the default values; otherwise, there may be insufficient buffers for the non-ISA-AA MDA, which may lead to packet discards. [117290]
- The use of AARP on multihomed, active-active SAPs or spoke-SDPs will force some of the traffic to use the inter-shelf AARP shunt interfaces. The AA remote divert will override policy-based routing (such as for NAT forwarding) applied on filters for traffic from the AARP instance (SAP or spoke-SDP).
- When **detect-seen-ip** is enabled in a **transit-ip-policy**, the operator must ensure that a default **app-profile** is configured. If there is no default **app-profile** and an **app-profile** is not provided by either RADIUS, Diameter or DHCP, then AA subscriber creation will fail; however, traffic for that subscriber will continue to traverse the AA on the parent context.

10.43 BFD

- **per-fp-egr-queuing** for LAG-based SAPs that have BFD sessions should not be enabled. When **per-fp-egr-queuing** is configured on a LAG and fast BFD is enabled for any SAP interface on that LAG, the BFD packets may be dropped on egress during LAG physical or logical port oversubscription. This condition may lead to the BFD session going down.

10.44 BFD on LSPs

- Interoperability with non-SR OS implementations of LSP BFD is not supported in Release 14.0.R4.

10.45 BFD VCCV

- The following table describes BFD VCCV interoperability with JunOS running on Juniper MX. [185090]

Table 28 **BFD VCCV Interoperability with Juniper MX**

Service	Interoperability
BGP-VPLS	BFD VCCV inter-op not supported
LDP-VPLS	BFD VCCV inter-op supported
Epipe control-word	BFD VCCV inter-op supported
Epipe no-control-word	Inter-op not supported
VPWS control-word	Inter-op not supported

10.46 BGP EVPN and XMPP Interoperability with Nuage

- In general, the recommended version to be used with Release 13.0.R4 is Nuage 3.2.R1 and higher for XMPP interoperability.
- The use of the “Policy-based Forwarding/Routing to an EVPN ESI” feature, for the integration of the SR OS nodes in the Nuage Service Chaining architecture, requires Release 3.2.R1 or higher in the Nuage VSC.
- The use of XMPP for the Fully-Dynamic VSD integration model requires Release 3.2.R1 or higher in the Nuage VSD. If lower VSD release versions are to be used, the following compatibility matrix provides an indication of the combinations that work or do not work:

Table 29 **Nuage VSD and SR OS Node XMPP Compatibility**

Nuage VSD Release	SR OS Release	Compatibility	Comments [1]
3.0.R3 – R5	12.0.R7 – R9	✓	S-D only
	13.0.R1 – R2	✓	S-D only
	12.0.R10 and higher	X	—
	13.0.R3 and higher	X	—
3.1	Any	X	Not a DC version
3.0.R6 and higher	12.0.R7 – R9	X	—
	13.0.R1 – R2	X	—
	12.0.R10 and higher	✓	S-D only
	13.0.R3 and higher	✓	S-D only
3.2.R1 and higher	12.0.R7 – R9	X	S-D only
	13.0.R1 – R2	X	S-D only
	12.0.R10 and higher	✓	S-D only
	13.0.R3	✓	S-D only
	13.0.R4 and higher	✓	S-D and F-D
	14.0.R1 and higher	✓	S-D and F-D
4.0 and higher	14.0.R1 and higher	✓	S-D and F-D

Note:

1. S-D = Static-Dynamic model, F-D = Fully-Dynamic model.

- A number of changes have been progressively introduced in the Nuage and SR OS EVPN-VXLAN implementation in order to align the control plane with the relevant IETF standards. In general, the use of SR OS Release 13.0.R4 and Nuage Release 3.2.R1 or higher is recommended. If lower release versions are to be used, the following compatibility matrix provides an indication of the combinations that work or do not work for EVPN. Note that if VSD – SR OS node integration is required, the above table must also be considered

Table 30 Nuage VSP and SR OS Node EVPN Compatibility

Nuage Release	SR OS Release	Compatibility	Comments
Up to 3.0.R3/3.1.R2	12.0.R7	✓	—
	12.0.R8/13.0.R*x*	X	Incompatible extended community values: RFC 5512 BGP encapsulation and Router's MAC.
3.0.R4-R6/3.1.R3	12.0.R7	X	Incompatible extended community values: RFC 5512 BGP encapsulation and Router's MAC.
	12.0.R8 and higher	✓	—
	13.0.R1 and higher	✓	—
3.2.R1/3.0.R8 and higher	12.0.R7-R8	X	Different VNI encoding can create issues
	13.0.R1	X	Different VNI encoding can create issues
	12.0.R9 and higher	✓	—
	13.0.R2 and higher	✓	—

- Notes: the following changes have been implemented along the releases:
  - The standard EVPN extended community values were introduced in Nuage Release 3.0.R4/3.1.R3 and SR OS Release 12.0.R8. Before those releases:

        - The VXLAN tunnel value in the RFC 5512 BGP encapsulation extended community was not compliant with *draft-ietf-bess-evpn-overlay.*
        - The Router's MAC extended community type/sub-type was not compliant with *draft-ietf-bess-evpn-prefix-advertisement*.
    - From SR OS Releases 12.0.R9/13.0.R2 and higher, the label field is interpreted as a 24-bit value when the encapsulation is VXLAN and it is ignored. Up to these releases, the SR OS node was expecting the Bottom of Stack (BoS) bit set in the label field.
    - From Nuage Release 3.2.R1 on, Nuage encodes the VNI in both, the Ethernet Tag and label fields. It can accept VNIs from both fields.
    - From SR OS Release 13.0.R4 on, the SR OS node encodes the VNI in the label field. It can accept VNIs from both fields.
    - Note that support for AD routes (EVPN route type 1) on the SR OS node has been introduced in SR OS Release 13.0.R4. Prior to that release, the SR OS node would discard any AD route received from VSC.
    - Nuage Release 3.0.R7 is not recommended in combined SR OS node and Nuage EVPN deployments.

10.47 BGP-EVPN Services

- Unknown unicast frames received on SAPs on an EVPN-MPLS enabled VPLS service use multicast-queues instead of unknown-queues. This should be taken into account when planning the QoS configuration.
- If both the following conditions are present:
    - **config**>**router**>**bgp disable-communities extended** is configured in a router with EVPN services
    - the service encapsulation does not match the configured **config**>**router**>**bgp def-recv-evpn-encap** encapsulation type (MPLS or VXLAN)

  then BGP-EVPN routes may need to be re-advertised after a CPM/CFM High-Availability switchover.

  For example, when **config>router>bgp disable-communities extended** is configured and if the router is configured with **def-recv-evpn-encap mpls**, EVPN-VXLAN services will have to re-advertise EVPN routes after a CPM/CFM switchover.
- When adding a new all-active Ethernet Segment (ES) on a node, use the following procedure to avoid potential transitory loops/black-holes for CEs in BGP-EVPN VPLS services:

1. Shut down the port corresponding to the ES in the PE (this will also ensure that the CE does not send traffic towards the PE while the ES is being configured).
2. Execute the **configure** and **no shutdown** commands on the ES.
3. Wait a few seconds for the exchange and process BGP-EVPN ES and AD routes to connect.
4. Execute the **no shutdown** command on the port.

Nokia also recommends that the configuration of a port **hold-time up** greater than zero on the ports associated to the ES. Upon a node recovery event (after reboot or node failure) the **hold-time up** value will give enough time to the core network protocols to setup the connectivity before allowing the CE to send traffic to the network. [214893]

- When PBB Source-BMAC is changed in a PBB-EVPN B-VPLS service, a **bgp-evpn mpls shutdown** or **bgp-evpn mpls no shutdown** is required for subsequent CMAC-Flush notification messages to use the latest PBB Source-BMAC (applicable to single-active Ethernet Segments using PBB Source-BMAC). [248860]
- In a scaled scenario, typically when a new BGP peer is added, there is a potential risk of having temporary **leaf-ac** to **leaf-ac** BUM traffic between EVPN E-Tree PEs. If the ingress PE receives the egress PE's Inclusive Multicast route prior to the leaf ESI-label, BUM frames from the **leaf-ac** will be forwarded to the egress PE without the leaf ESI-label, preventing the egress PE from filtering egress traffic to **leaf-ac**s. The filtering will work as soon as the egress PE's leaf ESI-label is received and programmed. [250969]
- If **route-target** family, **mp-bgp-keep**, or a local EVPN service are not configured prior to the ISSU (In-Service Software Upgrade) to Release 15.0.R4, EVPN routes will not be automatically advertised by an ABR/ASBR following the upgrade. Without **route-target** family, **mp-bgp-keep**, or a local EVPN service, after an ISSU upgrade, the BGP peer needs to be bounced to trigger EVPN route advertisements.

10.48 PBB-EVPN E-Tree

- An I-VPLS E-Tree service should not be linked to a non-EVPN B-VPLS service. Although the CLI will allow this association, Nokia recommends avoiding this association unless it is done for migration purposes.
- **pbb**>**leaf-source-bmac** is not restricted when configured along with I-VPLS E-Tree and B-VPLS services without BGP-EVPN.
- If an I-VPLS E-Tree service is used in a non-EVPN B-VPLS, leaf AC traffic will be sent to the B-VPLS network with a BMAC SA = **leaf-source-bmac**.

- Just as two given PEs cannot be configured with the same **source-bmac** so that traffic is not dropped, two PBB-EVPN E-Tree PEs cannot be configured with leaf-source-BMACs that match other leaf-source-BMACs or source-BMACs in the network.

10.49 E-Tree

- In EVPN E-Tree, ETH-CFM MACs for MEPs on SAPs and SDP bindings are always advertised as root MACs, irrespective of the access circuit being a leaf or a root. Therefore, unicast CFM-generated tests between MEPs on two remote **leaf-ac** instances will not be filtered as expected.

11 Known Limitations

The following sections describe the known limitations for SR OS Release 15.1.R3.

Note:

- Bracketed [ ] references are internal tracking numbers.
- Known limitations added in this release are marked **[NEW]**.

11.1 Hardware

- The AUX port on the SF/CPM or CFM is not supported in software. SR OS does not provide a means of configuring the device.
- The SyncE/IEEE 1588 port on the CCM-X20, CPM5, CPM-a, and CPM-e are not supported (reserved for future use).
- The LCD panel on the CCM-X20 is not supported (reserved for future use).
- The E-SATA interface on the CPM-X20 is not supported (reserved for future use).
- The link LED and operational status of a 10GBASE WAN-PHY port is tied to the Ethernet channel's ability to obtain frame-lock, so if there is a SONET issue such as PPLM, the link LED will not be lit, even though the SONET connection might otherwise be valid. [35354]
- A SONET/SDH port that is shut down or in internal loopback is incorrectly being allowed as a valid synchronous timing reference. [36448]
- The 3HE04116AA (SFP – 100/1000 FX SGMII 2KM ROHS 6/6) functions as dual-rate only when used with another 3HE04116AA. [67690]
- After a CFM High-Availability switchover with a c8-chds1, c4-ds3 or c1-choc3-ces-sfp CMA, if the system detects a configuration mismatch between the CFM and CMA, the CMA will automatically reset and the following message will be displayed on the console (for example, on MDA slot 1): “redDynamic:WDDI:winpathHwAudit Configuration out of sync between SF/CFM and MDA 1. Clearing the MDA to recover.”. [67797]
- When an m1-choc3-ces-sfp or m4-choc3-ces-sfp MDA is installed in an IOM3-XP/-B/-C, a larger-than-expected phase transition may be experienced when performing an adaptive clock recovery. [78408]
- A limit of two MDAs of type ATM, ASAP or CES are supported in a 7750 SR-c4/c12 system. For example, the limitation is reached with one m4-atmoc12/3-sfp and one m12-chds3-as. This applies to MDAs only and not to CES CMAs.

- On the 7750 SR-c4/c12, the 5-port GigE CMA cannot co-exist beside any of the other lower-bandwidth CMAs (including 1-port GigE and other lower-speed interfaces) in odd-even slot pairs (for example, slots 1 and 2, 3 and 4, 5 and 6, 7 and 8, 9 and 10 and 11 and 12). However, it is possible to have a 5-port GigE CMA in slot 2 beside a 1-port GigE in slot 3.
- Due to event suppression of Ethernet port states, a port that bounces while transitioning up or down may not take on its steady state for at least a second. Any port hold-timer configuration of less than one second will effectively look like a one second hold-timer. [91563]
- When the active and inactive CPM types are different, the provisioned card-type for both the active and inactive CPM will display the card-type of the active CPM. The equipped card-type will still display properly. [105862]
- When a differential DS1 on a CEM CMA/MDA is deleted and reconfigured as a differential E1, the recovered clock on the E1 may go into holdover. The clock recovery can be restored on the E1 with the CMA/MDA **clear** command. [109738]
- On the m4-chds3-as and m12-chds3-as MDAs, when a DS1 channel with SF framing and no occupied timeslots is active, the remote port will interpret its content as containing an RAI signal. This cannot be prevented, but only occurs when there are no channel-groups configured on the channel. If there are one or more channel-groups configured on the channel, it will still intermittently send “phantom” RAIs. However, this can be prevented by configuring at least one group to have “idle-cycle-flags ones”. This issue does not affect other ASAP MDAs. [129991]
- For 802.3 clause 50 compliant operation of 10G WAN-PHY ports on either SONET or SDH infrastructure, only the use of the SONET (default) framing option is supported (that is, **config>port** *port-id***>sonet-sdh>framing>sonet**). Although the system allows the user to configure **framing sdh**, this is an invalid configuration on a 10G WAN port. Interoperability issues may occur when attempting to use any of the following card types in SDH mode: m1-10gb-xp-xfp, m2-10gb-xp-xfp, m4-10gb-xp-xfp, imm4-10gb-xfp, imm8-10gb-xfp, imm5-10gb-xfp, and icm2-10gb-xp-xfp. [131400]
- On the 10GE HS-MDAv2 when the **agg-rate-limit** option is enabled for subscribers in a subscriber-profile, strict priority scheduling among traffic classes is not always maintained. To achieve strict priority scheduling, use subscriber **agg-rate-limit** in combination with **port-scheduler-policy** or **exp-secondary-shaper**. [159449]
- The 1-port 10GE HS-MDAv2 FPGA has a per-queue limit of around 2 Gb/s at a 64 byte fixed frame size. For a frame size of 64 bytes, the user needs at least five HS-MDAv2 queues for the full 10 Gb/s port bandwidth with 2 Gb/s per queue. For higher frame sizes (around 400 bytes), full 10 Gb/s can be achieved with a single queue. [166778]

- When a 10G DWDM tuneable SFP+ (3HE08142BA) reports signal-failure, the port will remain up. [211495]
- The Ethernet port hold timers are not synchronized across the redundant 7750 SR-c12 CFMs. In case of hold-time up, a dual CFM switchover within the period of the specified hold time will result in port state flaps for ports that have such a hold time configured. The workaround is to avoid performing a dual CFM switchover within a period of time that is lower than any of the configured hold times. [237356]
- When BGP on the router advertises FlowSpec routes to EBGP peers, non-transitive extended communities are not stripped from the advertised routes.
- When PW-Port Ether-type is set to a non-default value (value other than 0x8100), the **vc-type** command under the PW-Port (non FPE based PW-Port) is disabled.
- For L2oGRE Termination on an FPE-based PW-port, the following limitations apply:
  - L2oGRE tunnels can be terminated only on a non-system IP address in the Base-routing context
  - **vlan-vc-tag**, **force-vlan-vc-forwarding**, or **force-qinq-vc-forwarding** commands under the **spoke-sdp gre-eth-bridged** CLI context are not supported
  - Only IPv4 GRE transport is supported
  - L2oGRE can be only terminated on a PW-port. No construct other than PW-port can be provisioned in a **vc-switching** Epipe that contains a L2oGRE spoke-SDP (of type **gre-eth-bridged**). For example, another spoke-SDP (of any type) or a regular SAP (1/1/1:10) cannot be configured in the same **vc-switching** Epipe that contains a L2-GRE spoke-SDP.
  - Dual-homing redundancy using MCS in ESM is not supported
  - SRRP is not supported
  - LI is supported on PW-SAP but not on L2oGRE SDP or PW-port
- Egress FCS-error alarms may in some cases report invalid source card slot numbers for slots that do not have a card equipped. [257024]

11.2 Satellites

- If a satellite is to be moved to a new host chassis, or to a set of uplinks previously associated with a different satellite on the same host chassis, or is to be deconfigured and reconfigured with a new satellite ID, it should be reset first with the **admin satellite eth-sat** *sat-id* **reboot** command.
- Ethernet half duplex is not supported on Ethernet satellite (7210 SAS-Sx) ports.

- Only fiber SFPs can be used with combo ports (ports 1 and 2) of the Ethernet satellite.
- The fixed copper port associated with combo ports (ports 1 and 2) are not currently supported.
- On the 7210 es64-10gb-sfpp+4-100gb-cfp4 satellite, when 10G ports are **shutdown**, they will not report a remote fault sent by the peer, even when configured to do so. All other cases where remote faults are generated are handled correctly. [251427]

11.3 System

- Port-level and SAP-level statistics do not reflect packets processed by the CPM or CFM, for example, packets destined to a router IP address or a packet with the router alert options set. Another case is where DHCP relay packets ingress on a spoke-SDP bound to an IES interface as these packets are first sent to the CPM or CFM, so the SDP does not reflect that these are ingressing packets. [16330]
- The 7750 SR-7/12/12e and 7450 ESS-7/12 chassis cannot differentiate between a missing and non-functioning fan tray. [17756]
- The CLI allows the user to specify a TFTP location for the destination for the **admin save** and **admin debug-save** commands which will overwrite any existing file of the specified name. [18554]
- Dropped incoming packets due to a packet processing error are not being counted in the ifInErrors SNMP counter. Examples of packets such as this include any packet with a malformed IP header. [27699]
- Collision events detected on a CPM or CFM management Ethernet port are reported as CRC/Alignment errors. [30205]
- All IOM/IMM/XCM-based statistics (port, interface, and so on) are locally maintained on the IOM/IMM/XCM, not the CPM. IOM/IMM/XCM counters are not cleared when a **clear** command is issued; the CPM stores the reference values for the last clear operation and calculates the new values based on the values reported by the IOM/IMM/XCM. The reference values are not maintained between the active and standby CPM, so if a CPM switchover occurs, the newly active CPM will display the current values read directly from the IOM/IMM/XCM regardless of any clear command issued on the other CPM. [30444]
- When a fan is removed from a 7750 SR-7/12/12e or 7450 ESS-7/12, an erroneous “fan high temperature” alarm is generated that is cleared when the fan is replaced. [36112]

- Source address configuration applies only to the Base-routing instance, and where applicable, to VPRN services. As such, source address configuration does not apply to unsolicited packets sent out the management interface.
- TIMETRA-PORT-MIB.mib does not include an entry for “Link Length support” as an attribute of a Gigabit Ethernet port. This prevents Nokia NFM-P (formerly 5620 SAM) from reporting the value even though this attribute is reported in the CLI. [46225]
- After 497 days, system up-time will wrap around due to the standard RFC 1213 MIB-II 32-bit limit. [137937, 200196]
- Remapping of control plane traffic from a default CPM queue to a different queue is not supported on the 7750 SR-c4/c12. [59438]
- When the **password aging** option is enabled, the reference time is the time of the last boot and not the current time. Password expiry will also be reset on every reboot. [64581]
- In-service upgrades from SF/CPM3 to SF/CPM4 and from SF/CPM3/4 to SFM5/CPM5 are not supported.
- PCS High BER conditions on Ethernet ports are not being alarmed as a separate alarm condition and are incorrectly reported as a Local Fault. [98366]
- The **no debug** command does not remove the debug mirror information. [115892]
- Although extracted control traffic that arrives on a network interface but inside a tunnel and logically terminates on a service is supposed to bypass the Distributed CPU Protection (DCP) function, VPRN trace packets (**oam vprn-trace**), in this case, will be subject to DCP.
- The following considerations apply to the IF-MIB enhancements:
  - The **enable-ingress-stats** option must be enabled in CLI in order to increment the ingress IF-MIB counters for transit traffic. Ingress IF-MIB counters are updated even if a packet is discarded on an incoming interface. ifInDiscards is incremented if a packet is dropped as a result of a uRPF failure.
  - If a drop filter is configured on an incoming interface, ifInDiscards counters will be updated for IES/VPRN interfaces, but not for Base router or **management** router interfaces.
  - The following commonalities exist between IES/VPRN and Base router or **management** router interface counters:
    - Discard packets that need fragmentation but the DF bit is set: ifOutDiscards is updated
    - Discarded Broadcast-traffic: InDiscard is not updated
    - Data traffic is not reflected in the counters for a tunnel interface. Only control traffic (for example, LDP, RSVP, OSPF, IS-IS) will update the counters for a tunnel interface

    - Multicast traffic is reported in the unicast counters, but will not be reported in the case of a tunnel interface.s
    - Counters in the ifXTable and ifTable of the IF-MIB may not be updated properly during a High-Availability switchover or after a **clear router interface statistics** command. [146878]
- Too many files in a single subdirectory can result in longer read or write operations and eventually cause performance degradation of applications that regular need to access the compact flash. This is a limitation of FAT file system. [192499]
- OOB management Ethernet port redundancy is not supported during boot-up. Both management IP addresses must be on the same IP subnet.
- Configuration rollback is not supported across major releases. The software release major version of a node on which a **rollback revert** is being executed must match the software release major version used to produce the rollback checkpoint.
- After executing the CLI command **tools perform system script-control script-policy stop all**, queued EHS scripts are not executed. [234444]
- The Quality Level advertised on synchronous Ethernet (SyncE) connections on the Extension chassis of a 7950 XRS-40 is the Quality Level of the master chassis. This means that the extension chassis must be traceable to the same source as is used by the master chassis. Refer to the *7950 XRS-20 and 7950 XRS-40 Chassis Installation Guide* for details on the proper installation cabling to facilitate this traceability.
- On a 7950 XRS-40, the Extension chassis **sync-if-timing** will wrongly report free run after activity switch on the Extension chassis when the Extension chassis is in holdover state. [252695]
- When a port on an me2-100gb-cfp4 or me2-100gb-qsfp28 MDA is used as synchronous Ethernet (SyncE) reference into the central clock and an LOS condition on this port is detected, the central clock will switch to another reference if available. During this switch a phase transient that exceeds the limit defined by the standards may be observed. [253138]
- IEEE 1588 Port-Based Timestamping (PBT) is not supported on ports of Ethernet satellites.
- PXC is not supported on ports in DWDM, WAN or OTN mode.
- PXC ports do not support:
    - Dynamic Port Buffer Allocation (Named pools)
    - **eth-tunnel**s and **eth-ring**s
    - 802.1x Authentication
    - MC-LAG

  - Micro BFD on a LAG with PXC member ports (Micro BFD is enabled directly in the LAG context where BFD executes directly on individual member ports).
- Log events appear in the log recording time (timestamp) in chronological order; however, minor logging slowdowns may cause some log recording times to appear out of order.
- When iom4-e-HS scales to 96K SAPs, these specific scenarios should be avoided:
  - Majority of the services are Epipe services
  - Both SAPs of each Epipe are on the same MDA
  - Epipe SAPs are all LAG SAPs

  If all of these conditions are true, a rapid **shutdown**/**no shutdown** of the MDA that all Epipe SAPs are residing on might cause transient system instability.
- The following features are not supported on an IOM4-e-HS:
  - Port cross-connects (PXC)
  - Ethernet satellite host ports
  - Reset card on recoverable error
- The following commands and branches are present in the CLI but are not supported:
  - **configure system management-interface configuration-mode**
  - **configure system management-interface cli cli-engine**
  - **configure system management-interface cli md-cli**
  - **configure system grpc auto-config-save**

11.4 RADIUS

- If the system IP address is not configured, RADIUS user-authentication will not be attempted for in-band RADIUS servers unless a source-address entry for RADIUS exists.
- The NAS IP-address selected is that of the management interface for out-of-band RADIUS servers. For in-band RADIUS servers if a source-address entry is configured, the source-address IP-address is used as the NAS IP address; otherwise, the IP-address of the system interface is used.
- SNMP access cannot be authorized for users by the RADIUS server. RADIUS can be used to authorize access to a user by FTP, console or both.

- If the first server in the list cannot find a user, the server will reject the authentication attempt. In this case, the router does not query the next server in the RADIUS server list and denies access. If multiple RADIUS servers are used, the software assumes they all have the same user database.
- In defining RADIUS Vendor-Specific Attributes (VSAs), the TiMetra-Default-Action parameter is required even if the TiMetra-Cmd VSA is not used. [13449]
- Configuring a **fallback-action** under **config>subscr-mgmt>authentication-policy** to **accept** should not be combined with managed SAPs. Instead, Nokia recommends setting **fallback-action** to **user-db** *name* and configuring a default host to catch all entries and to provide default values for managed-SAP parameters.

11.5 TACACS+

- If the TACACS+ **start-stop** option is enabled for accounting, every command will result in two commands in the accounting log.
- If TACACS+ is first in the authentication order and a TACACS+ server is reachable, the user will be authenticated for access. If the user is authenticated, the user can access the console and any rights assigned to the default TACACS+ authenticated user template (**config**>**system**>**security**>**user-template tacplus_default**). Unlike RADIUS, TACACS+ does not have fine granularity for authorization to detail if the user has just console or FTP access, but a default template is supported for all TACACS+ authenticated users.

  If TACACS+ is first in the authentication order and the TACACS+ server is NOT reachable, authorization for console access for the user is checked against the user's local or RADIUS profile if configured. If the user is not authorized in the local/RADIUS profile, the user is not allowed to access the node.

  Note that inconsistencies can arise depending upon combinations of the local, RADIUS and TACACS+ configuration. For example, if the local profile restricts the user to only FTP access, the authentication order is TACACS+ before local. If the TACACS+ server is UP and the TACACS+ default user template allows console access, an authenticated TACACS+ user will be able to log into the console using the default user template because TACACS+ does NOT provide granularity in terms of granting FTP or console access. If the TACACS+ server is DOWN, the user will be denied access to the console as the local profile only authorizes FTP access. [39392]

11.6 CLI

- Non-printable, 7-bit ASCII characters are not allowed inside the various description fields. [93998]
- Output modifiers (“**| match**” and “**>**”) are not supported in configuration files executed using the **exec** command (scripts).
- Candidate commands (for example, **candidate edit**) cannot be used in an **exec** script and cannot be used in a cron job.
- A candidate configuration (created via **candidate edit**) is not preserved when a CPM/CFM failover occurs (the candidate will be empty).

11.7 Ingress Multicast Path Management

- The **show mcast-management channel** command does not show counts of the replications on the ancillary path. [65824]
- Multicast traffic may be affected for 10 seconds on a Soft Reset of the ingress card. [76417]
- Ingress multicast traffic through a queue with multipoint-shared queuing enabled will not be managed by IMPM when IMPM is enabled on the same ingress complex. [82402]
- Individual MMRP group entries cannot be displayed via CLI. [84252]
- When multicast traffic is received over a multicast tunnel using RFC 6037 MVPNs with all channels de-encapsulated from the multicast tunnel and terminating on the local PE with Ingress Multicast Path Management enabled on the related ingress FP, then the **show mcast-management channel** and **tools dump mcast-path-mgr channels** output may display a small amount of bandwidth for the channel corresponding to the multicast tunnel. This is expected and occurs due to the difference in the measured bandwidth of the channels between subsequent polls.

11.8 DS1/E1

- Via SNMP, a value of zero will be returned for tmnxDS1BERTTotalBits as this function is not supported on the DS1/E1 CMA. This value is properly shown as “N/A” in the CLI. [bz1400]

11.9 SONET/SDH

- On the m16-oc12/3-sfp, m8-oc12/3-sfp, m16-oc3-sfp, m8-oc3-sfp, m4-atmoc12/3-sfp, and m16-atmoc3-sfp MDAs and the c2-oc12/3-sfp CMA, LOP-P defects received by the MDA/CMA are incorrectly reported as AIS-P events. [8658]
- The **show port** command on a SONET/SDH interface will only display the bottom 4 bits of the S1 byte but will incorrectly display the bits as an entire byte. [17364]
- CV errors are incorrectly being incremented during a Severely Errored Seconds (SES) state. [29052]
- On the m1-oc192, m4-oc48-sfp and m2-oc48-sfp MDAs, if the H1 and H2 bytes are set to 0xFF but the H3 byte is not set to 0xFF, an AIS-P condition is not reported but an LOP-P condition is reported. [30498]
- The system does not prevent the user from entering more than 15 bytes in a path trace field for ports that have been configured for SDH framing; however, the system will only use the first 15 bytes of the entry for the path trace. [99733]
- OC-12c/STM-4c, and OC-48c/STM-16c and OC-192c/STM-64c SONET/SDH interfaces only run in CRC32 mode. CRC16 mode cannot be configured for these interfaces.
- On the m16-oc12/3-sfp, m8-oc12/3-sfp, m16-oc3-sfp, m8-oc3-sfp, m4-atmoc12/3-sfp, and m16-atmoc3-sfp MDAs and the c2-oc12/3-sfp CMA, only the first 16 bytes of the 62 byte trace string can be unique for each group of four ports (for example, for ports 1 through 4 or 13 through 16) for ports operating in SONET mode at OC-3. The last 48 bytes of the trace string will be the same for all ports and will be the last value set. Basically, a unique trace string per port is not possible if the unique part of the string is longer than 14 characters.
- On the m16-oc12/3-sfp, m8-oc12/3-sfp, m16-oc3-sfp, m8-oc3-sfp, m4-atmoc12/3-sfp, and m16-atmoc3-sfp MDAs and the c2-oc12/3-sfp CMA, the normal range for the SONET/SDH line signal failure Bit Error Rate (BER) threshold configured using the **config port** *port-id* **sonet-sdh threshold** command is 3 to 6. For these MDAs and CMA, the allowed threshold values are 3 to 5. The SNMP variable for this exponential threshold is tmnxSonetBerSfThreshold.
- The ports on the m16-oc12/3-sfp, m8-oc12/3-sfp, m16-oc3-sfp, m8-oc3-sfp, m4-atmoc12/3-sfp, and m16-atmoc3-sfp MDAs and the c2-oc12/3-sfp CMA are serviced in groups of four (1-4, 5-8, 9-12, 13-16) by a single framer chip, and as such, all must have the same framing across a given group. If framing on one port is changed, all four ports in a group must be **shutdown** and the framing will be changed on all four ports.

- The framer on the m4-oc48-sfp and m2-oc48-sfp MDAs supports a single software reset for all transmit subsystems, so changes to the transmit clock source on a single port will result in a short traffic interruption on all ports on the MDA. As a result, a short interruption will be experienced on all ports on the MDA when the transmit clock source for any one port is changed, for example from line to node timed. Also, traffic will be interrupted on all ports on the MDA when the port loopback mode on a port also configured with loop timing are transitioned in any of the following ways:
  - from “no loopback” to Internal
  - from Internal to “no loopback”
  - from Internal to Line
  - from Line to Internal.
- Receiving an LOF-E1 error condition on an E1 channel on the c1-choc3-ces-sfp CMA will cause the system to incorrectly raise an RAI alarm in addition to the expected OOF alarm on that E1 channel. [114221]
- On the m4-oc48-sfp-b, m16-atmoc3-sfp-b, m4-atmoc12/3-sfp-b and m16-oc12/3-sfp-b MDAs, a change to the transmit clock source on a port will result in a short interruption on that port. [119314]

11.10 Frame Relay

- If several MLFR links are removed rapidly from a bundle, one of the links may be deleted before it can send a remove-link message. If this occurs, the far-end link will not be notified and traffic loss may be seen until the far-end link times out and becomes non-operational. This will not occur if the DS0 group or the T1/E1 interfaces are shut down first, or if the links are removed a few seconds apart. [75883]

11.11 TDM

- When a TDM channel is administratively disabled, the alarm statuses from **show port** are correct; however, the alarm log “Alarm RAI Set” is only reported when the condition is cleared. [58505]

11.12 PPP

- PPP is not preventing IPCP negotiation with a non-matching IP subnet address. [24475]
- For MLPPP network port bundles and bundle-protection groups, PPP keepalive traffic is shown in the egress network queue statistics, but not in the egress port statistics.

11.13 ATM

- ATM ports whose operational state toggle at a high rate (faster than both the up and down hold timers) may remain in a “Link Up” but not be in the operationally Up state. The workaround is to wait for the hold timer to expire before issuing the **no shutdown** command. [35066]
- ATM port statistics for AAL5 packets include all AAL type frames as well as ATM cells received on L2 ATM pseudowires (Apipes) on the OC-3c/STM-1c and OC-12c/STM-4c ATM MDAs. This does not apply to an ASAP MDA. [39089]
- If the receive side fiber of an ATM Apipe SAP loses link and that Apipe is also bound to an SDP, then remote OAM cells received on that SDP will be dropped since the Apipe service is locally in a down state. Additionally, ETE-RDI cells will be transmitted out the ATM SAP to the CE. [39571]
- On the OC-3c/STM-1c and OC-12c/STM-4c ATM MDAs (and not an ASAP MDA), ATM Apipes configured with **vc-type atm-vpc** drop all ATM OAM F4 segment cells and pass through the ATM OAM F4 end-to-end cells. The PTI field of the forwarded ATM OAM F4 end-to-end cells is set to five and might cause interoperability issues if the third-party equipment expects the PTI field to be zero. [40451]
- Bi-directional FR PVC management procedures over an ATM VC part of an FRF.5 VLL are not supported. When doing FRF.5 interworking between different models of SR/ESS or other products, the bi-directional network PVC management over the ATM VC must be disabled on the other products. [49696]
- If traffic is passing on an ATM OC-12 port and the port speed is changed to OC-3, “Unknown Protocol Discards” may be seen at the console although no such frames are actually being received. The OC-3 port's operational state is not affected, although some noise may be interpreted as end-to-end VC-RDI/AIS cells by newly configured ATM PVCs, which would cause those PVCs to go operationally down. The condition will clear as soon as ATM traffic passes once again through the port. [58197]

- ATM cells in a VPC connection with the GFC field not equal to zero will be discarded. This only affects non-ASAP ATM MDAs. [75387]
- See SONET/SDH in the Known Limitations section for additional limitations that affect ATM MDAs.
- On the OC-3c/STM-1c and OC-12c/STM-4c ATM MDAs (not the ASAP MDAs), some ingress traffic counters do not update for certain types of ATM OAM F5 cells. This results in discrepancies between the ingress traffic statistics: PVC vs. port vs. SAP, packets vs. octets. Egress traffic is not affected. [109427]

11.14 ATM MDAs Access Mode Only

- The ATM interfaces on non-ASAP MDAs listed below only support the customer-facing access mode.

Table 31 **ATM MDAs that Support Access Mode Only**

Nokia Part #	Description
3HE00074AA	16-port ATM OC-3c/STM-1c MDA – SFP
3HE00071AA	4-port ATM OC-12c/STM-4c MDA – SFP
3HE05944AA	16-port ATM OC-3c/STM-1c MDA – SFP Rev B
3HE05945AA	4-port ATM OC-12c/STM-4c MDA – SFP Rev B

See ASAP for more information about the ASAP MDA.

11.15 ATM and IS-IS

- IS-IS is not supported on IES and VPRN interfaces with ATM PVC SAPs in this software release.

11.16 ATM Traffic Management/ Statistics Limitations

The following limitations only apply to the OC-3c/STM-1c and OC-12c/STM-4c ATM MDAs and do not apply to the ASAP MDAs.

- In the context of multiple services using an ATM MDA, the following two criteria must be met in order to satisfy the QoS guarantees:
  - VC fairness
  - COS fairness
- VC fairness implies that each VC gets its due share of bandwidth relative to the other VCs and COS fairness implies that within each VC, each COS gets its due share of bandwidth. What is considered the "due share" is very specific to the configuration. (For example, for two VCs of the same ATM service category, the due share will be proportionate to the configured rates of the VCs; for two VCs with different ATM service categories, the due share will depend on the priority of the service category and the configured rate, and so on.)
- A minor loss of throughput (< 2% of line rate) may occur if an OC-12 port is configured with small number of shaped PVCs, the difference in the configured ATM rates of the PVCs is large, and the sum of the shaped rates is equal to port rate. The loss of packet throughput occurs in the highest traffic parameter VC and only. [28869]
- The ATM layer shaping in the MDA schedules cells of the high-priority Forwarding Class queues with strict priority over cells of low-priority Forwarding Class queues within a SAP. This is performed such that packet delay and jitter are minimized on the high-priority forwarding class queues. As a result in some traffic loading scenarios, the lower priority forwarding class queues may not achieve their fair share of bandwidth. This is the case when the high-priority Forwarding Class queues have an offered traffic to the ATM MDA per-VC queue equal or higher than the PIR of the ATM VC. The user can alter this behavior and trade delay performance for forwarding class fairness in this specific scenario configuring H-QoS schedulers to limit the total offered load out of the forwarding class queues to the ATM MDA per-VC queue to the PIR of the ATM VC. [30819]
- OC-12/STM-4 latency increases when applying a new ingress SAP policy that adds more queues. The latency increases from around 22.2 *m*s to 24.8 *m*s over a 1 min period. Traffic loss does not occur during this period.
- Port input statistics do not increase when terminating e-t-e AIS cells are received.
- PVC admin state is not applicable. There is no command that can administratively disable a PVC; in order to disable a PVC, the user must disable the applicable service or service interface.

11.17 Class of Service Fairness Affected on Shaped VCs

- The following limitation only applies to the OC-3c/STM-1c and OC-12c/STM-4c ATM MDAs, and do not apply to the ASAP MDAs.

  In the case of ATM VCs configured with more than two classes of service where one queue, queue A, is allowed no burst beyond CIR and another queue of the same priority, queue B, is allowed to burst up to line-rate; the traffic offered to queue B might prevent queue A from achieving its CIR. The problem has a lesser degree of impact if there is an increased number of ATM VCs on the port and can also be addressed by lowering the configured PIR of queue B. [35224]

11.18 ASAP

- Following is a list of limitations for the 4/12-port Channelized DS3 MDA, the 1-port Channelized OC-12/STM-4 (DS0) and the 4-port Channelized OC-3/STM-1 (DS0) ASAP MDA:
  - BERT pattern 2e20 is not supported.
  - ATM ILMI support is not enabled.
  - IPv6 is supported for network mode PPP channels and access mode PPP, FR and cHDLC channels and MLPPP bundles.
- In exceptional cases, especially in a fully loaded node, where the occurrence of a High-Availability CPM or CFM switchover is exactly concurrent with an APS switch from Working to Protect (both unidirectional or bi-directional failures), PSBF may potentially be posted by the far-end node during the APS K1/K2 byte exchange due to the increase latency response of the near-end where the CPM or CFM switchover is occurring. [41192]
- DS3 configuration with m23 framing on the channelized ASAP MDA may detect false AIS. This may cause the DS3 to bounce occasionally. [74671]

11.19 LAG

- A failure of the link holding the primary port of the LAG can sometimes very briefly impact (<10e-4 seconds) flows on other links of the same LAG. This is not the case for failures on other links (non-primary) of a LAG. [49698]

- When **lag-per-link-hash** or **lag-link-map-profile** is used for a given SAP or network interface egress traffic, sub-second OAM traffic generated by the router (if supported for a given service/network interface) may not follow the same link as the data path traffic.
- When **lag-per-link-hash** or **lag-link-map-profile** is used for a given SAP or network interface egress traffic and BFD is enabled on that interface, BFD packets remain round-robin over the active links of the LAG irrespective of which link is used on egress by the given SAP/network interface.
- On a LAG, CPM-originated sub-second CFM/BFD packets use hashing independent of that configured for the data traffic. When **per-fp-egr-queuing** is enabled, the CFM/BFD packets may egress LAG over a different port than used by the SAP's data traffic. For those CFM/BFD packets, internal system queues, instead of the SAP's queues are used, and CFM/BFD packets are not accounted for in the SAP queues.
- Pulling out the active CPM/CFM can, in rare cases, result in LACP to signal to adjacent nodes that ports are going down. To avoid this and other potential issues, Nokia strongly recommends always pressing the RESET button before pulling out a CPM/CFM card. [146453]
- Access-egress queue optimization feature **per-fp-egr-queuing** is not supported on the same LAG with BFD. However, this restriction is not enforced. If BFD is erroneously enabled, BFD packets may use a different LAG port than the egress LAG port used for data traffic, and if the port is oversubscribed, the BFD packets may starve and lead to the BFD session going down. [155303]
- When BFD is to be originated/terminated in a SAP context on a given LAG with **per-fp-sap-instance** enabled, Nokia recommends using, at minimum, a one-second interval timer. Very large SAP scales on LAG may require even larger timer values, especially on older SR OS system. Failure to do so may result in BFD sessions going operationally down during LAG-member-port status changes. [170148]
- Multicast CAC supports up to eight levels per LAG; thus, the operator cannot define different levels for every possible LAG port count when LAG contains more than eight member ports. [175567]
- PW-SAP on distributed mode LAG with Vport is not supported. [178343]
- For mixed-speed LAG member port support, ingress-rate and egress-rate for LAG member ports must be set to default.

11.20 VSM

- The VSM-CCA-XP only provides ifInUcastPkts, ifInOctets, ifOutUcastPkts and ifOutOctets counters. The VSM-CCA-XP does not distinguish between unicast, multicast and broadcast packets. As a result, IP multicast statistics are also not supported on a VSM-CCA-XP IP interface. [40551]

11.21 MLPPP

- If several PPP member links in a MLPPP bundle are removed or shut down at the channel-group level simultaneously, term-requests may not be sent out. In this event, the far-end links may not be notified and the links may not become non-operational until PPP keep-alives fail. To work around this issue, shut down member links at the physical level first (if possible), or remove links or shut down channel groups one at a time. [87044]
- IPv6 interfaces over MLPPP bundles are only supported on ASAP MDAs even though the system allows that configuration on other MDA/CMA types. [143700]

11.22 APS

- Ports that are part of an MLFR bundle or that contain an MFLR bundle cannot be APS protected.
- APS is not supported on MDAs/CMAs that support LAN and WAN-PHY mode for 10G ports (for example, m2-10gb-xp-xfp).
- The imm1-oc768-tun card does not support APS.
- When an APS group contains circuits on separate ATM MDAs, both MDAs must be in the same ATM mode (max8k-vc|max16k-vc).
- Annex B (of ITU.T G.841) is supported in the following scenarios:
  - Supported with single chassis APS (SC-APS) only (no MC-APS support)
  - Supported on all 7750 SR/ and 7450 ESS platforms and with all IOM types.
- A mirror/LI destination SAP cannot be on an APS protected port.
- Restrictions specific to SC-APS:
  - Bundles are not supported on ports (or contain ports) that are protected with *uni-directional* SC-APS.
  - Uni-1plus1 SC-APS is supported only on the 7750 SR-c4/c12 platforms. Only the following cases are supported:

    - POS ports on non-channelized MDAs configured in network mode
    - CES ports configured in access mode where only Cpipe services (SAPs) are configured on that port.
    - ASAP channelized ports with MLPPP where the ports are configured in network mode.
- Restrictions specific to MC-APS:
    - Network mode ports cannot be part of an MC-APS group.
    - Ipipe SAP cannot be on a port that is part of an MC-APS group.
    - Routing protocols cannot be run over MC-APS protected ports (however, static routing is allowed).
    - BFD and VRRP over MC-APS protected ports are not supported.
    - The only type of bundle that can be *bi-directional* MC-APS protected is MLPPP with IPCP encapsulation (on ports configured in access mode).
    - Ports with Frame Relay (FR) or Cisco HDLC encapsulation cannot be protected with MC-APS.
    - Only *bi-directional* mode is supported with MC-APS. The *uni-directional* and *uni-1plus1* modes are not supported.
- In some cases of RDI-L, the transmitted K1/K2 bytes on the wire may differ from those maintained by the CPM or CFM's APS controller (as displayed in CLI). [36537]

11.23 TCP Authentication Extension

- It is not possible to delete an authentication keychain if that keychain was recently removed from a BGP neighbor while BGP was operationally down. BGP has to become operationally active before the keychain can be deleted. [57277]

11.24 SNMP Infrastructure

- After an SNMP log is removed and recreated, traps will no longer be sent to a **trap-target** that has the **replay** option configured. To start sending traps again, the **trap-target** should be removed and recreated. [162559]

11.25 Routing

- Setting a metric of zero in OSPF or IS-IS is not supported and causes the interface to fall back to the **reference-bandwidth** computed value instead of setting the value to zero. [17488]
- Routes exported from one protocol to another are redistributed with only the first ECMP next-hop. Therefore, if BGP routes having multiple next-hops are exported to a VPRN client, only one next-hop for the route will be exported. The one chosen is the lowest IP address of the next-hop address list. [40147]
- A static route with a CPE connectivity target IP address which is part of the subnet of the static route itself will not come up if there is no alternate route available in the routing table which resolves the target IP address. This is because a static route can only be activated if the linked CPE session is up, and in this case the CPE session can only come up if the static route itself is activated. [62663]
- Policy-statement entry **from interface** *interface***-***name* can only be used with multicast routing and will not match other routing protocols. To achieve a similar match for other routing protocols, **from protocol direct** with a prefix-list should be used. [89371]
- When the applied export policy is changed in conjunction with an **export-limit**, it may not take effect immediately without clearing the policy (**no export/ export**), or in very few cases, toggling the administrative state of the protocol. [90244]
- There is no warning trap sent after a clear export policy is issued when the **export-limit** is increased a few times and **clear export** is performed. [90274]
- Using **no preference** in the routing policy does not trigger re-evaluation of routes that are being leaked from another local VRF. The workaround is to set the preference with the desired value in the policy. [114322]
- Static routes do not take an IPv6 Anycast address as next-hop. [115800]
- The LFA next-hop may use the same egress interface as the primary next-hop when a mix of IES spoke-SDP interfaces and network interfaces is present. [141276]
- uRPF and interface statistics may not be correct after an event such as a **clear statistics**, **clear card** or **switchover**. [150500]
- If the **triggered-policy** command is enabled, in order for route policies to take effect after a High-Availability switchover, **clear** commands must be executed or the **triggered-policy** configuration toggled (**shutdown/no shutdown**). [154937]
- IP options 131 (Loose Source and Record Route) and 137 (Strict Source and Record Route) are not processed. Destination-based routing will be performed on the IP packets containing these options. [167864]

- A **clear** of the uRPF statistics should only be performed when uRPF is enabled for IPv4 and IPv6. If not, the counters may not reset to zero. [174961]
- NG-MVPN inter-AS Routing Options B and C for multicast has the following known limitations :
  - No source can be supported at the ASBR routers.
  - For MLDP in GRT and NG-MVPN option C, Basic Recursive Opaque MLDP FECs are used as per RFC 6512 section 2.
  - For NG-MVPN option B, where the system IP of the root node is not visible to the leaf nodes in the non local ASs, Recursive Opaque MLDP FECs are used as per RFC 6512 section 3.
  - This feature is only supported for dynamic MLDP.
- Configuring an **arp-timeout** of less than a minute could result in unexpected ARP-table refresh behavior and traffic impact. [226590]
- The **tools perform service id** *service-id* **interface** *ip-int-name* **ignore-sap-port-state** command is accepted and can become active on a satellite Ethernet port SAP IP interface. However, the IP interface will remain non-responsive. [234262]
- The following limitations apply to secondary IP address scales on access interfaces.
  - The IPsec, GRE, L2TPv3, and IP-in-IP protocols are supported at the old scale of 16 secondary IP addresses. Configurations are not to exceed 16 secondary IP addresses when these protocols are active on the interface.
  - The current number of 16 secondary IP addresses still applies on network interfaces.

11.26 IP/RTM

- The traffic sent to non-subsuming routes of an aggregate route with an indirect next-hop address to be resolved by a VPN-leaked route will be black-holed. [149804]
- Routes are flapped for a static-route (indirect) which is resolved via IS-IS when an LFA change occurs, even though the primary next-hop for IS-IS does not change. [251403]

11.27 Routing Policies

- In a routing policy configuration that exports routes into IS-IS, the statement **to level** sets the level of the route and is not a match criteria. However, if an incompatible level is specified or the destination protocol is not IS-IS, then no match is returned and policy evaluation stops. For example, if the router is configured as L1 only and **to level 2** is specified, then policy evaluation stops and will not evaluate subsequent entries.

  On the router redistributing the BGP routes into IS-IS, an IS-IS export policy containing two entries is applied. The first entry matches, except for the **to level 2** statement because the router is configured as L1 only. The second entry is a full match. Both entries have an **action accept** statement, so the BGP-learned routes should be redistributed into IS-IS (by entry 2). However, due to the behavior outlined above, this does not happen and no routes are exported from BGP into IS-IS.

  To avoid this condition, the correct IS-IS level should be set or the statement should be omitted. Alternatively, an entry with a **to level** statement should be placed at the end of a policy. [171345]
- Policies using the action **next-entry** do not operate as expected when the following condition is true: a route-policy statement with two entries, for which some routes match the first entry but not the second one. If the action in the first entry is **next-entry**, the action of the second entry will be irrelevant since the routes do not match. One might expect that the routes would be processed as configured in the default action of the policy. However, they will behave as the default action of the protocol to which the policy is applied. [173046]

11.28 IPv6

- When **debug router ip packet** is enabled, packets received on a 6-over-4 tunnel do not display the IPv4 header information and packets sent on the tunnel do not display the IPv6 header information as the encapsulation and decapsulation is performed on the line card. [45606]
- The following restrictions apply for IPv6 support for HTTP redirect:
  - no support for ESM Wholesale/Retail
  - no support for one-time HTTP redirect
  - no support for ESM credit-control IPv6 filters
  - ingress only

11.29 DHCP

- If the addition of the Option 82 information to a DHCP packet would cause the maximum size of 1500 bytes to be exceeded, the DHCP relay incorrectly does not forward the original DHCP packet (without the additional Option 82 information). [37061]
- A Local User Database (LUDB) cannot be applied to the DHCPv6 Local Server used for ESM.
- In Releases 11.0.R1 and higher, PPPoX leases are no longer persistent (stored on compact flash) in an SR OS-based DHCPv4 server. [148366]
- A DHCP server using per-pool **failover** is not allowed to synchronize with a DHCP server using per-server **failover**. [169222]
- A DHCPv6 server in SR OS only accepts relayed messages (Relay-forward).
- DHCPv6 Relay-Forward messages received on an IPv6 interface that connects to a DHCPv6 client will be delivered to the DHCPv6 server in the following scenarios:
  - a single Lightweight DHCPv6 Relay Agent (LDRA) in front of an ESM subscriber interface
  - DHCPv6 Relay Agents in front of an ESM subscriber interface with DHCPv6 snooping enabled at the group interface. A combination of LDRA and DHCPv6 Relay Agents is supported with a maximum of five.

  The following examples are not supported:
  - an LDRA in front of a regular (non-ESM) IPv6 interface
  - a DHCPv6 Relay Agent in front of a regular (non-ESM) IPv6 interface
  - a DHCPv6 Relay Agent in front of an ESM subscriber interface with DHCPv6 snooping disabled at the group interface
- A forceRenew message from a DHCP server, located in the same VRF as the DHCP relay, is sent as a unicast message to the client’s IP address. The source address of the forceRenew is the actual DHCP server IP address while it should be the one configured as **siaddr-override** address. [212028]
- If both nodes’ MCS databases are in synchronization, a **no shutdown** of the **local-dhcp-server** with **failover** enabled could result in one side getting in Normal state, while the remote side stays in pre-Normal state for MCLT time before moving in Normal state. [239195]

11.30 RIP

- The RIP global statistics for all RIP instances is incorrectly being displayed for each VPRN instance. This has the effect of causing one to think that the VPRN instance has learned routes when in fact it has not. [26472]
- When 16 bytes of **authentication-key** was configured in RIP, the last byte was filled with the null character in Release 10.0 and Release 11.0 prior to 11.0.R6. Interoperability issues would arise when the network consisted of SR OS routers running these older releases and those running 11.0.R6 or higher. [167905]

11.31 IS-IS

- ECMP across multiple-instances is not supported. ECMP is per instance only. Only one route, the one with the lowest instance ID, is installed. [85326]
- In a multi-instance IS-IS configuration, the same IS-IS prefix is not leaked to all instances with Level-1 and Level-2 leaking. Leaking between instances is configured with routing policies. [85463]
- There is no separate **export-limit** configuration for IPv6 in IS-IS. The same **export-limit** is used for IPv4 and IPv6 routes depending on the policy configuration. [91520]
- IP Fast Reroute (FRR) does not guarantee low loss when multiple interfaces are going down; it is limited to first-order failures where loop-free forwarding as a property continues to hold. It is possible that the loss is low because all down events are detected before the first IGP SPF runs, and, the updated topology does not result in a loop. Nokia recommends against depending on FRR in such topologies.

  SR OS defaults to one next-hop only in ECMP scenarios. In cases where ECMP paths exist, it is possible that the IGP chooses an Loop Free Alternative (LFA) that is different from any of the ECMP paths. While the FRR switch itself is (nearly) hitless, the subsequent IGP SPF-based next-hop update will pick one of the remaining ECMP paths as the primary next-hop. A change in the primary next-hop that is not the same as the previously computed LFA can result in transient forwarding loops, based on the updated topology. This could be especially amplified if the SPF timers are different, or if the routers in the network are heterogeneous (different vendors, different route processor speeds/ capability).

  Note that the same sequence of convergence events can occur, even if ECMP > 1 is configured, as long as there are more than MaxECMP paths available; the next-hop count of one is a special case of the same. [130305]

- When the LFA next-hop for a far-end GRE tunnel is activated, packets of a spoke-interface do not benefit from IP FRR but wait until the SPF has updated the new primary next-hop for the GRE SDP far-end before resuming forwarding. [130913]
- IP FRR degrades to regular convergence when IS-IS is the DR on a broadcast interface and the failure is a interface shutdown. As such, Nokia recommends a P2P configuration. [138279]
- In a network with a VPRN PE node redistributing BGP-VPN routes into IS-IS and an IS-IS level-1/2-capable CE router in the connected IS-IS network leaking these routes from level-1 to level-2 could result in a routing loop when the PE receives the level-2 route and replaces the BGP-VPN route with it so that it is no longer exported. A workaround is to tag all BGP-VPN routes that are exported to IS-IS and to block all tagged IS-IS routes from getting redistributed in level-2 on all level-1/2-capable CE nodes. [168803]
- When IGP-shortcut is not enabled for all prefix families (for example, it is enabled for IPv6 family and disabled for IPv4 family), a prefix of the family which is disabled and which inherits an LFA backup of type tunnel from a node in the SPF tree will not use the LFA backup and will remain unprotected. [251402]
- IGP-shortcut in an IS-IS instance does not support the use of an SR-TE LSP as an LFA backup next-hop. When enabling the **lfa-protect** or the **lfa-only** option in an SR-TE LSP configuration, a warning is issued. The IGP-shortcut feature can use an SR-TE LSP as a primary next-hop of an IPv4 prefix, an IPv6 prefix, or an LDP IPv4 prefix FEC. [261897]

11.32 OSPF

- The system may refresh self-originated LSA shortly after completing a CPM or CFM switchover which may mean the entry is refreshed before the expiration of the age-out period. [65195]
- An SR OS router with more than one point-to-point adjacency to another router over links of equal metric, may compute the shortest-path tree over the incorrect link in the case of unidirectional link failures on the far-end router. This condition lasts until the dead timer expires and the adjacency over the broken link is brought down locally (near-end). A workaround is to change to broadcast interfaces or enable BFD over these links. [79495]
- During High-Availability switchover, more than the configured **export-limit** routes get leaked when exporting to OSPF. Once the High-Availability switchover is completed, routes will come back as restricted by export-limit. [90098]

- The export limit will not show the export-count after route summarization; it only displays the routes exported before summarization. If the routes have not been advertised due to an OSPF **external-db-overflow** condition, the **export-limit** count will still count the routes as exported. [91520]
- When export limit is reduced via the **export-limit** command, toggling the administrative state of the protocol is required to remove all previously exported routes. [91520]
- IGP-shortcut in an OSPF instance does not support the use of an SR-TE LSP as an LFA backup next-hop. When enabling the **lfa-protect** or the **lfa-only** option in an SR-TE LSP configuration, a warning is issued. The IGP-shortcut feature can use an SR-TE LSP as a primary next-hop of an IPv4 prefix, an IPv6 prefix, or an LDP IPv4 prefix FEC. [261897]

11.33 OSPF PE-CE

- OSPF traffic engineering is not supported in VPRN instances.

11.34 BGP

- If BGP transitions to the operationally disabled state, the **clear router bgp protocol** command will not clear this state. The BGP protocol administrative state must be **shutdown/no shutdown** to clear this condition. [12074]
- If a 6PE prefix is received with two or more labels for the same next-hop, the reference count in the **show router bgp next-hop** output will always display a value of one. [56638]
- The system does not prevent the user from using the same IP address of a BGP peer on one of the router interfaces and configuring this can result in a configuration that fails to execute after a reboot. [57198]
- If the BGP neighbor address is configured prior to configuring that same IP address on a router interface, the configuration can be saved and loads properly with a warning message displayed. Also, the peering shows up as idle. The workaround is to not use the same IP address for a local router interface and a BGP neighbor. [85198, 132818]
- In a typical PE-CE scenario, when the PE is learning IPv6 routes from multiple CEs over a BGPv4 session, the traffic switchover time for IPv6 with EDGE-PIC may not be sub-100ms. To achieve this, a BGPv6 session protected by BFDv6 may be required to learn IPv6 prefixes. [122822]

- The BGP best route selected may change after two High-Availability switchovers when the **ignore-router-id** option is configured in the **bgp best-path-selection** context. [130406]
- When **local-as** is configured at the peer/group level, a set/reset of **local-as** at a higher level may cause the BGP session to flap. When **peer-as** is configured on the peer level, a set/reset **peer-as** on the group level will cause the BGP session to flap. [148704]
- If filter policy resources are not available for newly auto-generated address prefixes when a BGP configuration changes, new address-prefixes will not be added to impacted match lists or filter policies as applicable. The operator must free resources and change the filter policy configuration, or the BGP configuration must be changed to recover from this failure.
- Inter-AS options B and C are not supported between a confederation's member ASes. [157071]
- For inter-AS option C, BGP-3107 routes are installed into unicast RTM (**rtable-u**). Unless routes are installed by some other means into multicast RTM (**rtable-m**), Option C will not build core MDTs; therefore, **rpf-table** should be configured to **rtable-u** or both.
- When **update-fault-tolerance** is disabled, in some cases where the length of the aggregator, aspath, as4_aggr, as4_path attribute is wrong, an invalid-update log event is generated. [157817]
- The **clear router bgp protocol** command cannot be used to trigger BGP graceful restart (GR). It will clear the BGP routes before entering the helper mode. The proper way to trigger GR is to use the **clear router bgp neighbor** *x.x.x.x* command. [159793]
- If an SR OS node has negotiated graceful restart (GR) notification with a BGP peer and it detects a hold-timer expiry event, it will incorrectly display "hold timer expiry" instead of "send notification" as a reason for entering the GR helper mode in the **debug router bgp graceful-restart** output log. [161274]
- When **update-fault-tolerance** is enabled and all attribute length fields are okay, the peer is brought down when the mpreach/mpunreach attribute cannot be correctly parsed. [161501]
- The "Last Modified" timestamp in the **show router bgp routes detail/hunt** output can have the wrong value after a dual CPM/CFM switch over. [188240]
- When **next-hop-resolution use-bgp-routes** is configured, if **shortcut-tunnel** is configured with **disallow-igp** option, BGP routes do not get resolved over another BGP route.
- When a labeled-unicast route is leaked into the unlabeled RIB, or vice versa, the following limitations apply:
  - Split horizon behavior controlled by the **split-horizon** command is not respected.

  - Prepending of the **local-as** associated with the session over which the route was received is not supported.
  - The route table cost to reach the next-hop of the route is not available in the destination RIB and therefore cannot be used by the BGP decision process or to update the value in MED or AIGP path attributes.
  - The “stale” state of the route (due to GR) is not shown in the destination RIB.
  - The imported route is never grouped with other BGP routes in the same deterministic MED group, even if the neighbor AS is the same.
- BGP dynamic peer does not support:
  - **damp-peer-oscillations**
  - **graceful-restart**
  - **authentication-key**
  - **auth-keychain** [210255]
- The command **error-handling update-fault-tolerance** must be used in nodes running Releases prior to 14.0.R4, 13.0.R11, or 12.0.R20 when interoperating with routers that support VXLAN IPv6 transport, otherwise the router running the earlier release will bring down the BGP peer session. For BGP routers not supporting either IPv6 next-hops or **error-handling update-fault-tolerance**, a workaround could be the use of Route-Target Constraints to restrict IPv6 service route-targets from peers that do not support IPv6 services. This assumes that the Route-Reflector does support IPv6 next-hops. [233504]
- When an export policy with the next-hop set to an IP address falls in the same subnet as an EBGP peer's IP address, the advertised BGP routes by this EBGP peer, or by EBGP peers having similar settings in that group, can have inconsistent next-hops. [235321]
- For BGP routes, traffic does not flow through the LFA path if the LFA is resolved over a IGP-shortcut tunnel. [251643]
- The BGP minimum route advertisement interval (MRAI) is a per-peer timer, not a per-peer per address-family timer. As a result, when a route is selected for rapid-update advertisement to a BGP peer, all other pending route updates for that peer are also sent immediately, even if they are not included in the scope of the **rapid-update** command.
- BGP optimal route reflection is not supported for routes containing IPv6 BGP next-hop addresses.
- By default (without an **advertise-label pop** policy action), a router cannot originate a /32 label-IPv4 BGP route for which it has an active static, OSPF, or IS-IS route if there is no operationally-up tunnel to the /32 prefix.

11.35 BGP-EVPN

- BGP-EVPN MPLS is only supported in regular **vpls** and **b-vpls** services. Other VPLS types, such as **i-vpls** or **m-vpls**, are not supported.
- The **proxy-arp/nd** functions are fully supported in EVPN-MPLS services, including on SAPs/SDP-bindings that are part of an **ethernet-segment**. However **proxy-arp/nd** are not supported on I-VPLS.
- When **debug router bgp update** is enabled and EVPN-MPLS routes are received, the label-1 value shown in the debug output will not match the value shown in the **show router bgp routes evpn**. The debug output shows the entire 24-bit values as received on the route and **show** commands display the value interpreted as Label or VNI based on the received RFC 5512 tunnel-encapsulation extended community.
- In general, no SR OS-generated control packets are sent out to EVPN destinations. The only exceptions are CFM traffic (from UP MEPs, MIPs, and vMEPs), Proxy-ARP/ND messages (confirm messages), and IGMP messages.
  - **eth-cfm** MEPs and MIPs on SAPs and SDP-bindings are supported within EVPN-MPLS and EVPN-VXLAN VPLS services. EVPN-MPLS also supports full service-level MEPs (vMEP) which include extraction on the EVPN-MPLS connection. EVPN-VXLAN support for vMEPs does not include extraction for VTEP connections.
- xSTP and M-VPLS services:
  - xSTP can be configured in **bgp-evpn** services. BPDUs are not sent over the EVPN bindings.
  - **bgp-evpn** is blocked in **m-vpls** services, however, a different **m-vpls** service can manage a SAP/spoke-SDP in a BGP-EVPN-enabled service.
- In **bgp-evpn**-enabled VPLS services, **mac-move** can be used in SAPs/SDP-bindings; however, the MACs being learned through BGP-EVPN will not be considered.
- **disable-learning** only works for data-plane-learned MAC addresses.
- The following features and commands are not supported in combination with **bgp-evpn mpls**:
  - **mac-protect**
  - **bgp-vpls**
  - **endpoint** and attributes
  - Subscriber management commands under service, SAP and SDP-binding interfaces
  - **mld**/**pim-snooping** and attributes
  - **vsd-domain**

  - BPDU-translation
  - L2PT-termination
  - MAC-pinning
  - **spb** configuration and attributes
- ESI PBF is not supported across VPLS services (i.e., the interface on which the steering takes place and EVPN VPLS interface must be in the same VPLS service).
- BUM traffic matching an IPv4/MAC ESI PBF filter for EVPN will be unicast forwarded to the VTEP:VNI resolved through PBF forwarding.
- When **provider-tunnel inclusive mldp** is enabled in an EVPN-MPLS VPLS or B-VPLS service, in combination with **root-and-leaf** and **bgp-evpn**>**ingress-repl-inc-mcast-advertisement**, the system will send an Inclusive Multicast Ethernet Tag (IMET) route with a composite tunnel type in the Provider Tunnel Attribute. In releases up to and including Release 13.0.R7, BGP peers receiving these IMET routes will reset their BGP session unless **configure**>**router**>**bgp**>**error-handling**>**update-fault-tolerance** is enabled.
- P2MP MLDP support, when **provider-tunnel inclusive no shutdown** is enabled in an EVPN-MPLS service, has the following caveats.
  - The same IMET-P2MP route cannot be imported into two services at the same time. If that is the case, only one service will join the MLDP tree.
  - In general, the P2MP provider-tunnels have the following limitations:
    - **mac-ping**, **mac-trace**, **mac-populate** with **flood** option, and **mac-purge** with **flood** option are not supported
    - **sdp-ping** and **sdp-mtu** are not supported with a P2MP spoke-SDP used as an I-PMSI in a VPLS context
    - **p2mp-lsp-ping**/**trace** are not supported
- When **bgp-evpn mpls** is enabled in Epipes, the following caveats must be considered:
  - Epipes with **bgp-evpn** cannot be associated to a B-VPLS service.
  - No BGP-MH is supported.
  - The use of spoke-SDPs along with **bgp-evpn mpls** does not support the configuration of **vc-switching** on the Epipe.
  - No **endpoints** are supported in Epipes with **bgp-evpn**.
  - No **bgp-vpws** or **spoke-sdp-fec** configurations are supported.
  - The **pw-template-binding** command will not be blocked in **bgp-evpn** Epipes, but it will not have any impact on the service.
  - **ignore-oper-down** is not supported in **bgp-evpn** Epipes.

  - When setting up an EVPN-VPWS between PE1 and PE2, if the remote **eth-tag** in PE2 does not match PE1's local **eth-tag**, the Epipe service will be operationally up in PE1 but not in PE2. In order to avoid PE1 sending traffic that will be discarded at the egress PE2, **eth-cfm** can be used.
- For P2MP MLDP support for BGP-EVPN, when static P2MP MLDP tunnels and dynamic P2MP MLDP tunnels used by BGP-EVPN co-exist on the same router, it is recommended for the static tunnels to use a tunnel-ID lower than 8193. If a tunnel-ID is statically configured with a value equal or greater than 8193, BGP-EVPN may attempt to use the same tunnel-ID for services with enabled provider-tunnel and fail to set up an MLDP tunnel.
- The following features are not supported on EVPN MPLS R-VPLS Services:
  - I-VPLS on the R-VPLS
  - IP Multicast traffic on R-VPLS with **bgp-evpn** enabled. Hence the following commands are blocked for EVPN-VXLAN and EVPN-MPLS:
    - **allow-ip-int-bind forward-ipv4-multicast-to-ip-int**
    - **allow-ip-int-bind igmp-snooping**
    - **igmp**/**pim-snooping no shutdown** with R-VPLS and **vxlan**/**bgp-evpn mpls**
- When two BGP instances are enabled on the same VPLS service, the following features are not supported:
  - SDP-bindings
  - R-VPLS, M-VPLS, I-VPLS, B-VPLS or Etree VPLS
  - Proxy-ARP/ND
  - BGP multihoming
- A router with two BGP instances in the same service will not detect any duplicate MAC existing on the EVPN-VXLAN and EVPN-MPLS networks.
- The command **incl-mcast-orig-ip** is not supported in B-VPLS services.
- The **unknown-mac-route** command will trigger the advertisement of the unknown MAC route, only in the **bgp-evpn vxlan** instance.
- According to RFC 7432, when more than two PEs are part of a single-active Ethernet Segment (ES), a remote PE detecting the unavailability of the DF PE is expected to flush all of the MACs associated with the ES and flood any unicast traffic destined to that ES. However, in the current release and in this scenario, the remote PE will spray the unicast traffic among all remaining PEs in the ES without flushing the MAC addresses associated with the ES. [209329]
- **oam mfib-ping** is not supported in BGP-EVPN enabled services.

- The command **config**>**service**>**system**>**bgp-evpn**# **ad-per-es-route-target evi-rt-set** is not supported for EVPN E-Tree services. When the command is configured on a router, the AD per-ES routes (with ESI=0) used for EVPN E-Tree services are always advertised with the service route-target and route-distinguisher, irrespective of the **ad-per-es-route-target** configuration. AD per-ES routes for non-zero ESIs (used for regular multi-homing) will be normally sent using either **evi-rt-set** or **evi-rt** based on the router's configuration.
- Although Conditional Static Black-hole MACs may be configured in a two BGP-instance service, they are not supported. [246324]
- The following services and features in the context of Ethernet Segment (ES) are not supported with **enable-inter-as-vpn** or **enable-rr-vpn-forwarding** commands:
  - auto-discovery per-ES based mass-withdrawal for EVPN-MPLS services when the ES PEs and the remote PE are in different ASs or IGP domains
  - EVPN multi-homing when the ES PEs are in different ASs or IGP domains, or there is an NH-RR peering the ES PEs and overriding the ES route next-hops
  - IGMP/PIM snooping on a PE that is a also an ABR/ASBR
- PBB-EVPN destinations cannot support MPLSoUDP tunnels. An attempt to resolve a B-VPLS BGP-EVPN next-hop to an MPLSoUDP tunnel will fail, and the **show router bgp next-hop evpn** command will show that the next-hop is not programmed with a reason "Label StackLimit".
- EVPN VPLS/Epipe and R-VPLS destinations can use MPLSoUDP tunnels, as long as:
  - No options that add extra bytes to the egress packets are configured. An example of these options is entropy-label.
  - No **configure**>**system**>**ip**>**allow-qinq-network-interface** is executed on the router.

  An attempt to use any of the above options, or configure **allow-qinq-network-interface**, will result in a failure to resolve the BGP-EVPN route's next-hop.
- BGP-EVPN Ethernet Segment (ES) routes need to be resolved by BGP before the router can use them for DF election. The next-hop can only be resolved by a regular route in the routing-table (for instance, a tunneled shortcut route would not resolve an ES route's next-hop).

11.36 BGP VPWS

- If a multihoming PE receives a BGP-VPWS NLRI with the D-bit set or the CSV set from a remote PE, it will not cause the BGP-MH site within the service to go operationally down (and will subsequently cause a BGP-MH DF switchover). An example of this is if the remote PE shuts down the SDP connected to the multihoming PE; this will not cause a DF switchover on the multihoming PE. In order to achieve a DF switchover in this case, some kind of continuity check between the two nodes will be required (for example, SDP keepalives). However, network failures that cause the network PW on the multihoming PE to go operationally down will cause a DF switchover. [147804]
- If a BGP update for a VPWS service is received with a Circuit Status Vector (CSV) length field of greater than 32 bits, it will be ignored and not reflected to BGP neighbors. If a BGP update for a VPWS service is received with a CSV length field of greater than 800 bits, a notification message will be sent and the BGP session will restart. BGP VPWS services support a single access circuit; consequently, only the most significant bit of the CSV is used on transmit. On receive, for designated forwarder selection purposes, only the most significant byte of the CSV is examined.

11.37 Segment Routing

- When the preference is set the same for different Segment Routing (SR) protocols, SR protocols are not picked as per the default TTM preference but as per "route owner value". Hence, SR-OSPF is preferred over SR-ISIS, which is preferred over SR-TE. [219330]
- In case of an SR-TE LSP with multiple hops configured in the path, if the adjacency-label changes in an intermediate hop after the LSP has come up, there is no way for the head-end to get this new label (without doing an LSP **shutdown** followed by **no shutdown**) in case of locally-computed SR-TE LSP. This is not an issue for a PCE-controlled path.

11.38 MPLS/RSVP

- The **no rsvp** command in the **config**>**router** context has no effect as the state of RSVP is tied to the MPLS instance. The **no mpls** command deletes both the MPLS and RSVP protocol instances. [8611]

- An invalid Class Number or C-Type in the Session Object does not cause a PATH Error message to be generated. [12748]
- To disable OSPF-TE on a link, both ends of the link should be MPLS/RSVP-disabled for CSPF to work correctly and be removed from the TE database. [15127]
- The **bandwidth** parameter is not supported on PATH and RESV messages of one-to-one detour and facility-bypass paths. [27394, 57847]
- For (rare) topologies in which the protected LSP and the detours are set up along parallel links across several hops (link protection only), Fast Reroute (FRR) may take longer to restore traffic if the primary path is broken. [39808]
- Shutting down a port on an OC-3c/STM-1c MDA may not provide sub-50 ms failover for an RSVP path signaled over that port. This issue does not occur if the fiber is disconnected or if the path is shut down. [39973]
- Fast failover times of less than 100 ms cannot be achieved for Fast Reroute (FRR) protected LSPs if the failed link is detected by copper Ethernet SFPs. Sub-second failover times are achieved, but the failover times with copper Ethernet SFPs are inherently longer based on how the system communicates with the SFP. [49003].
- A manual-bypass tunnel that terminates on the incoming interface IP address at the merge point will become operational but will not be properly associated with the primary LSP. The recommendation is to always use the IP address of the system interface to ensure reachability to the node. [59184]
- 7750 SR-c4/c12 RSVP LSPs cannot be signaled over a channelized DS1 or E1 interface if the channel group bandwidth is less than 1 Mbps. [59776]
- There are scenarios where the bypass optimization does not ensure that a node-protect manual bypass will be selected over a node-protect dynamic bypass tunnel. This is because the manual bypass may be unavailable when the association of a bypass LSP is made with the primary LSP.

  The bypass optimization feature only changes the association for an LSP which requested node protection but is currently associated with a link-protect bypass.

  To ensure this selection when using manual bypass, dynamic bypass must explicitly be disabled. [60261]
- If a local IP address is configured with the same address as the destination address of an MPLS LSP, the LSP will no longer be set up and will use the RSVP error code of “routingError”. [73326]
- Least-fill behavior is not exhibited when the user does a configuration change MBB by decreasing the bandwidth on the LSP. [74544]
- In case of a non-CSPF LSP with only secondary paths, once the active secondary path goes down, the LSP will wait for the regular retry time. It will then try to set up again, and if that fails with a path error, it will go into fast-retry mode. [80012]

- On the leaf node of a P2MP LSP, the DSCP value of an IP packet will not be used for classification even though the **ler-use-dscp** option is configured in the network policy. The LSP EXP from the MPLS header will be used instead. The workaround is to not configure the **ler-use-dscp** flag on the network policy. [80105]
- Refresh reduction over inter-area manual bypass will only work if the RESV RRO format at the bypass destination is one of the following: IL, SLIL, SLI or SIL. [108420]
- For an LSP terminating or passing through a router where the OSPF router ID is different than the system interface, the AR hop table entry will be incorrect. [109589]
- If route recording is not enabled on manual bypass or the system interface is not recorded in RRO manual bypass, association of inter-area manual bypass to protected LSP may not work correctly. There may be an incorrect AR hop table entry when the OSPF router ID is different from system interface. Inter-area manual bypass association does work correctly for the following supported RESV RRO formats for the primary LSP path: SLIL, ILSL, SIL, SLI, ISL and SL.
  - S: RRO object with system ID
  - I: RRO object with interface ID
  - L: RRO label object

  If no node supports any of the formats above, the bypass LSP association to protect LSP may be incorrect. [109753]
- A manual bypass LSP may not come up if the user specifies a local interface address of a node in the **exclude-node** configuration of that LSP. When computing the CSPF path at the ingress (LER) or transit LSR (ABR), if the local interface is down or not part of the IGP or not in the same area as the node doing the CSPF computation, MPLS will be unable to resolve the interface address to its router ID and CSPF may not compute a path excluding the node specified by the user. [118046]
- MPLS-TP is only supported on static LSPs and static PWs.
- MPLS-TP LSPs can only carry static MPLS-TP PWs, while MPLS-TP PWs can be carried on static MPLS-TP LSPs or dynamic RSVP-TE LSPs.
- CAC is not supported for MPLS-TP LSPs or PWs.
- SVC-Ping and SDP-ping are not supported on MPLS-TP LSPs and PWs.
- Dynamic bypass LSP re-optimization does not support inter-area bypass LSP and P2MP LSP.
- Inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP are not supported.
- GMPLS LSPs are only supported on 10GE and 100GE ports.

- Penalty weights have no impact on backup LSP paths that are forced to be strictly SRLG diverse from the primary. That would be the case of secondary LSP paths and bypass backup LSP with the **srlg-frr strict** option enabled. When SRLG groups are changed on an MPLS interface on a node, this information is reflected on all other nodes, which have TE enabled and on which the IGP is not in administratively down state. Depending on the number of SRLG groups added or removed from an MPLS Interface, the expected results may not be immediately visible if SRLG groups are changed on-the-fly.
- An inter-area RSVP LSP with Fast Reroute (FRR) enabled or disabled but with the PATH message not containing the RRO may fail at an ABR with a failure code of “routingLoop”.
- A pre-empting LSR will perform hard pre-emption, instead of soft pre-emption if the PATH message of an LSP did not include the RRO.
- LSP BFD cannot be configured on RSVP LSP secondary paths.
- A CPM/CFM switchover will cause MPLS static LSP to flap which will have traffic impact on the users of the LSP.
- After a network churn, as IGP has completed converging, non-CSPF RSVP-LSP metric can be different compared with the relevant LSP metric contained in TTM. [220454]
- The following limitations apply to entropy labels:
  - Rolling back MPLS **entropy-label rsvp-te** to **force-disable** will fail if all of the following actions have been performed:
    1. Entropy label is disabled under RSVP, MPLS, LSP and in services.
    2. A **rollback save** is performed.
    3. Entropy Label is enabled under RSVP, MPLS, LSP and in services.
    4. A rollback is performed. [226474]
- MPLS Entropy Labels (RFC 6790) are not supported with L2TP and GTP tunnels.
- LSP BFD is not supported on LDP LSPs for LDP-over-RSVP. [245954]
- Entropy label is not supported for PCE controlled SR-TE LSPs.
- Generalized multi-protocol label switching (GMPLS) UNI is not supported on an IOM4-e-HS.

11.39 MPLS-TP

- **static-dynamic** pseudowire switching for MPLS-TP is only supported when the dynamic PW segment is a spoke-SDP using the PW ID FEC.

11.40 LDP

- If **triggered-policy** is configured, LDP policies are not dynamically evaluated for changes in FECs. [71830]
- It is not possible to apply an accounting policy in the egress LDP statistics context if both **default** and **record combined-ldp-lsp-egress** are configured in that policy. [84406]
- When enabling or disabling the **ldp-shortcut** option in the global routing context, any indirect LDP static-route will be operationally toggled and its age will be reset. [85366]
- A GRE SDP will stay operationally down in case the SDP far-end address resolves through an LDP or RSVP tunnel due to configured shortcuts. GRE tunnels cannot be established over MPLS tunnels. [92314]
- **clear router ldp instance** is not an atomic operation — it consists of **shutdown** followed by **no shutdown**. If a High-Availability switchover happens right after the **clear** command, the **no shutdown** part of the command might have been lost during the switchover, resulting in the LDP instance remaining shut down on the newly active CPM/CFM. After the switchover, the user can issue a **no shutdown** on the LDP instance to re-enable LDP. [160940]
- “Local Neighbor Liveness Time” and “Local Recovery Time” will not be updated in the existing session when a change is made to **graceful-restart maximum-recovery-time** or **neighbor-liveness-time**. Any new sessions established after GR timers have changed will use the changed values. [169756]
- The **ldp-sync** option can be enabled on a static-route entry in order to delay its activation in the event of an LDP discovery flap on the selected static-route next-hop interface.

  In the above scenario, if a CPM/CFM High Availability switchover occurs, the running **ldp-sync** timer could be incorrectly decremented, inducing an early activation of the static-route. [224939]

11.41 LDP IPv6

- The PW switching feature is not supported with LDP IPv6 control plane. As a result, the CLI will not allow the user to enable the **vc-switching** option whenever one or both spoke-SDPs use an SDP which has either **far-end** or **tunnel-far-end** configured as an IPv6 address.

- Layer-2 services that use the BGP control plane (such as dynamic MS-PW, BGP-AD VPLS, BGP-VPLS, BGP-VPWS, and EVPN MPLS) cannot bind to an IPv6 LDP LSP because a BGP session to a BGP IPv6 peer will not support advertising an IPv6 next-hop for the Layer-2 NLRI. These services will not auto-generate SDPs using LDP IPv6 FEC. In addition, they will skip any provisioned SDP with either **far-end** or **tunnel-far-end** configured to an IPv6 address SDP when the **use-provisioned-sdp** option is enabled.
- Multihoming with T-LDP active/standby FEC 128 spoke-SDP using LDP IPv6 LSP to a VPLS/B-VPLS instance is supported. BGP multihoming is not supported because BGP IPv6 does not support signaling an IPv6 next-hop for the L2 NLRI. The Shortest Path Bridging (SPB) features will work with spoke-SDPs bound to an SDP which uses an LDP IPv6 FEC.
- The following LDP capabilities are not supported with LDP IPv6:
  - resolution of IPv6 FEC or IPv4 FEC over a RSVP IPv4 LSP, where the FEC has been signaled by a IPv6 T-LDP session
  - resolution of IPv6 FEC over a RSVP IPv4 LSP, where the FEC has been signaled by a IPv4 T-LDP session

11.42 IP Multicast and MVPN

- The Router Alert IP option is not included in **mtrace** queries that are unicast to the last-hop router in the trace as defined by the IETF draft. Note that this causes no known interoperability issues since this packet is still destined for an IP address on this last-hop router. [37923]
- (S,G) or (*,G) multicast streams transmitted through an LAG will no longer be hashed on the UDP source or destination ports; identical streams with differing UDP ports will all transit over the same link. [66618]
- When a multicast CAC (MCAC) policy is applied under IGMP-snooping of a SAP with static-groups that are configured in the bundle of the same MCAC policy, the bandwidth used by the static groups on the SAP is not recalculated after the bundle is disabled and re-enabled. The used bandwidth remains at zero for the static groups. In addition, the MCAC recalculation command **tools perform service id** *service-id* **mcac sap** *sap-id* **recalc policy** *policy-name* fails to recalculate the used bandwidth, and the use of the **bundle** option in the command returns an error. [71023]
- When MoFRR for PIM is enabled, tunnel interfaces (for example, dynamic in-band MLDP interfaces) are ignored for MoFRR functionality.
- Some multicast limits (for example, the number of OIFs per IIF per line card) are not enforced by the system; thus, Nokia recommends that operators verify with Nokia support teams that planned deployment limits are supported.

- RPF Vector must be enabled on every router for RFC 6037 MVPN inter-AS option B/C. Failure to do so will result in RPF Vector being dropped and result in PIM Join/Prune processing as if RPF Vector was not present.
- Packets arriving on the standby interface that belong to a standby stream for a given (S,G) will be discarded and counted as either discards or mismatch against the (S,G) record. If the standby interface and the RP interface are identical, then a discard counter is incremented. If the standby interface differs from the RP interface or the RP interface is NULL, then a mismatch counter is incremented.
- MoFRR active joins are untouched when periodic **mc-ecmp-balance** rebalancing is active to prevent traffic impact.
- Deploying the sender-only/receiver-only feature requires all PE nodes in an ng-MVPN using RSVP P-tunnels to use SR OS Release 11.0.R1 or higher. [154000]
- When dynamic MLDP signaling is deployed, a change in Route Distinguisher (RD) in the root node is not acted upon for any PIM (S,G)s on the root node until the leaf nodes learn about the new RD (via BGP) and send explicit delete and create with the new RD.
- Enhanced multicast load-balancing (**config>system>load-balancing>mc-enh-load-balancing**) is mutually exclusive with PIM LAG usage optimization (**config>router>pim>lag-usage-optimization**), since CPM-based load-balancing cannot mimic data-path-based load-balancing in general cases (source IP unknown). Enabling both options at the same time is not blocked, but may lead to multicast traffic disruptions and thus, must be avoided. [179614]
- Packets arriving on the standby interface that belong to a standby stream for a given (S,G) will be discarded and counted as either discards or mismatch against the (S,G) record. If the standby interface and the RP interface are identical, then a discard counter is incremented. If the standby interface differs from RP interface or RP interface is NULL, then a mismatch counter is incremented. Auto-rebalancing when a new path becomes available is performed for active joins.
- When multicast source geo-redundancy is enabled, MCAC may incorrectly account for suppressed joins; therefore, Nokia recommends against enabling MCAC together with the multicast source geo-redundancy feature. [185533]
- For NG-MVPN inter-AS Routing Options B and C, configuring static MLDP and dynamic MLDP on the leaf could result in unexpected behavior if the LSPs' P2MP identifiers overlap.
- For NG-MVPN inter-AS Routing options B and C, configuring static MLDP on LEAF1 AS1 and dynamic MLDP on LEAF2 AS2 could result in unexpected behavior on the ASBR if the LSPs' P2MP identifiers overlap. In this case, the ASBR can merge the LSPs into a single uplink LSP toward the ROOT node. As such, the same multicast stream may be incorrectly flooded to both static and dynamic MLDP LSPs.

11.43 IGMP Reporter

- IGMP reporter has the following limitations:
  - no support for MLD (IPv6 multicast)
  - only supported on subscriber interfaces
  - no SAM support as collector device (collector device, in general, is not a part of IGMP reporter)
  - fixed MTU of 1400 bytes

11.44 PIM

- In certain VPLS topologies where multiple multicast sources are connected to different PEs configured with VPLS services using PIM-snooping, traffic duplication can occur on the egress SAP/SDP. This is due to the PIM-snooping/proxy with (S,G)/(*,G) interaction not working in accordance with *draft-ietf-l2vpn-vpls-pim-snooping-06* (Appendix B.2). [125379]
- In dual-homing PE scenarios where the path from the active source-PE to customer RP fails and recovers, a customer's channel (S,G) entry may remain programmed on the PE's VRF even if the receiver leaves the group. [152632]
- Nokia recommends using a minimum of 3.5 seconds hold time (Hello Interval times Hello Multiplier) on PIM interfaces and to use BFD if faster link-failure detection is required. [171934]

11.45 PPPoE

- HTTP redirect is not supported for L2TP sessions at the LAC. Attempting to use HTTP redirect IP-filters in ESM SLA-profiles that would be applied to L2TP sessions will block the HTTP traffic on those sessions. [81316]
- L2TP tunnel over GRE spoke-SDPs on an interface in a VRF is not supported.
- When configuring **reject-disabled-ncp** below the PPP policy, the system will only reply to a "PPP LCP Protocol Reject" message when an IPv6CP request is received while IPv6 is not supported. An IPCP(v4) request while IPv4 is not supported will still be silently discarded. [115620]

- With an incomplete SRRP setup for PPPoE subscriber hosts, IPv6 traffic originating on the backup node of an SRRP pair may be sent towards the subscriber host if SRRP was not active, causing that traffic to be dropped at the client. [117550]
- Host-tracking Multi-Chassis Synchronization (MCS) is not supported on PPPoE hosts.
- To support L2TP, UDP port 49151 is used for internal communication. Care must be taken this port is not blocked by any cpm-filter entry. [143110]
- For active PPPoE sessions in a dual-homed setup with DHCP leases granted via the internal DHCPv4 client and DHCP server, care must be taken when shutting down SRRP or taking it into an INIT state on both sides of the dual-homed setup. This will no longer result in a timeout of the PPPoE sessions but the granted lease can still time out on the DHCP server. The DHCP server then offering the same IP address to another DHCP client can result in a conflict: “PPPoE session failure on SAP *sap-id* in service *svc-id* - … PPPoE session with same IP * already exists in service *svc-id*”. To avoid these conflicts, either a shutdown of the related group or subscriber interfaces or a manual clearing of the hanging PPPoE sessions on both sides of the dual-homed setup must be executed. [203892]
- With **new-qinq-untagged-sap** disabled, the oldest PPPoE session can be terminated due to an LCP echo timeout when both single- and double-tagged PPPoE sessions are active on a SAP with QinQ encapsulation :*X*.0 (where *X* is any VID value different from zero (0)). Enabling **new-qinq-untagged-sap** prevents double-tagged sessions to become active on a SAP with QinQ encapsulation :*X*.0. A separate SAP must be created for double-tagged PPPoE sessions in this case. [234099]

11.46 QoS

- In a SAP ingress QoS policy with shared queuing, high-priority packets dropped will be counted in the low-priority drops of the SAP ingress service queue statistics. [32335]
- When provisioning a network port on an MDA results in more than 8192 ingress queues needing to be allocated on the MDA, the CPM and IOM can show different usage numbers for ingress queues in certain situations. When this happens, the numbers will synchronize back up when the newly-provisioned network port is deconfigured. [32878]

- When **ler-use-dscp** is enabled on network ingress and multicast VPRN traffic is tunneled through an SDP, ingress classification on network ingress will happen based on the TOS bits in the transport (outer) IP header as opposed to the customer IP packet. This behavior is seen strictly in multicast VPRN packets. [40348]
- When the router is operationally down in a VPRN instance because the route-distinguisher is not yet defined and PIM is then enabled on a VPRN SAP, the CPM will allocate multicast queues for the SAP whereas the line card will not allocate queues because the line card does not know that multicast is enabled on the interface. This disparity in allocation of queues will exist only in the transitional phase until the route-distinguisher is set after which the line card will allocate multicast queues and the line card and CPM will be in synchronization. [42469]
- Network control traffic (or other high-priority, expedited traffic) should not be configured to share a queue on a port scheduler policy with non-expedited or lower priority traffic or the queue could get into a state where the higher priority traffic will not be forwarded out the egress port. This can also occur if the traffic is on two separate queues that are mapped to the same level. [59298, 59435]
- Small amounts of packet loss may occur on queues configured with an MBS equal to or lower than 4 KB and/or lower than two times the maximum packet size of packets forwarded by these queues. This can happen when the traffic rate through these queues is large or when there is a large amount of jitter on this traffic. This packet loss is possible on queues where the traffic rate is lower than the PIR. To avoid this type of packet loss, the MBS of a queue should be configured to a minimum value of 5 KB or to two times the maximum expected packet size, whichever is higher. [66687]
- When sizing the mega pool based on the buffer-allocation requirements, the size is rounded up to the nearest available value and may result in no buffers being available for other pools. In non-named-pool mode, all port pools are guaranteed a minimum size of 16k (which is rounded up to 6 buffers=18k). This guarantee does not apply to **named-pool-mode** and named pools still have no minimum size (could be zero), but MDA default pools now have a minimum size of 1 Mbyte. [80716]
- When the **agg-rate-limit** option is enabled on a Vport used by a subscriber, any subscriber host queue that is parented to a virtual scheduler is not rate-limited by the Vport aggregate rate. The queue will compete for bandwidth directly on the port's port scheduler, at the priority level and weighted scheduler group at which the virtual scheduler is port-parented. If the virtual scheduler is not port-parented, or if there is no port scheduler policy on the port, the host queue will be orphaned and will compete for bandwidth directly based on its own PIR and CIR parameters. [109318]

- WRR distribution across CVLANs will not be correct for certain combinations of **class-agg-weight** and frame size, such that frame size**/class-agg-weight** results in a value lower than 64 bytes. The system will round up the value resulting from frame size**/class-agg-weight** to be at least 64 bytes. A few examples of such combinations are: 200-byte frames and weight 8, 100-byte frames and weight 4, and 70-byte frames and weight 2. [112010]
- Network egress queue-groups cannot be used for frames coming from the CPM or CFM other than IPv4, IPv6 and MPLS types. Other frame types (for example, ARP or IS-IS) egress out of the per-port network-queue mapped to FC NC instead of the queue-group queue. [115427]
- The advanced-config-policy **sample-interval** H-QoS parameter is supported only for policers and not for queues. [125417]
- In-profile broadcast, unknown unicast and multicast traffic that is accounted as offered-combined by a multi-point service queue is accounted as offered-uncolored in the forwarding engine statistics on FP3-based line cards. [128123]
- Out-of-profile unicast traffic that is accounted as offered-colored by a unicast service queue is accounted as offered-hi-priority in the forwarding engine statistics on FP3-based line cards. [128133]
- When applying an ingress network-queue policy on an MDA that belongs to an IOM with only one complex (that is, IOM3-XP/-B/-C) or that is inserted in a 7750 SR-c4/c12 chassis, the network-queue policy will also be applied to the other MDAs belonging to the same IOM or the same chassis. [138995]
- When **enqueue-on-pir zero** is enabled on a queue, the PIR of the queue is not set to zero immediately for inactive queues. Instead, the setting is applied only after the queue’s next scheduling opportunity.
- The combination of Ethernet tunnels configured with access LAG emulation **adapt-qos** distribute mode and an egress port scheduler is not supported. Since a port can be a member of more than one **eth-tunnel** and those **eth-tunnel**s could have different **adapt-qos** modes, anything at the port level (like **port-scheduler-policy**, port queue-groups queues, port queue-group schedulers and arbiter, **agg-rates**) will be unaffected by the **eth-tunnel adapt-qos** mode.
- The **port-fair** mode on **eth-tunnel** will calculate the rates based on the number of active paths and not based on the path bandwidth.
- When the CBS and MBS for a queue have similar or equal values, the system automatically changes the CBS value to be larger than configured. This ensures that a request for a buffer from the reserved pool is honored correctly when there are available buffers in the reserved part of the queue’s pool. This does not change the operation of the MBS, which continues to be the maximum drop tail for the queue. [149831]
- 802.3 SNAP frames are supported on SAP ingress QoS classification as part of MAC criteria. IP QoS reclassification works only for Ethernet II or PPPoE frames at SAP egress; it does not work with 802.3 SNAP frames. [188450]

- On egress, IPv4 QoS-based classification criteria are ignored when MAC-based ACLs are configured.
- Concurrent MAC-based QoS/filter policy match criteria and IPv6-based QoS/ filter policy match criteria are not supported on access interfaces. On ingress, IPv6 routed packets ignore MAC-based QoS classification criteria, while switched packets ignore IPv6-based ACL match criteria. On egress, IPv6 QoS-based classification criteria are ignored when MAC-based ACLs are configured. [208461]
- If an automatic data-path recovery action occurs on the 7750 SR-a4/a8, causing a control-protocol failure, it is possible that no tmnxEqDataPathFailureProtImpact alarm is raised. [209067]
- When a SAP egress QoS policy is applied to a B-VPLS SAP, any classification using **ip-criteria** or **ipv6-criteria** statements is ignored for PBB-encapsulated traffic; the classification does apply to non-PBB traffic egressing the B-VPLS SAP.
- When a SAP ingress QoS policy is applied to a B-VPLS SAP, any classification using **ip-criteria** or **ipv6-criteria** statements will apply to PBB-encapsulated traffic except in the case of IPv6 traffic when two inner VLAN tags are present.
- Remarking of the inner dot1p or DE bits based on the profile result of egress policing is not supported.
- Self-Generated Traffic Quality of Service (**sgt-qos**) for Diameter only marks traffic to the well-known destination port 3868. If a different port is configured in the Diameter peer policy (**configure aaa diameter-peer-policy** *peer-policy-name* **peer** *name* **transport tcp port** *port*), then the **sgt-qos** configuration for the Diameter application does not become active.
- Egress-policed packets can be directed to a local SAP queue, and, when this is configured, the output of a **show service id** *service-id* **sap** *sap-id* **sap-stats** only counts these packets through the policer; that is, they are not counted a second time through the queue to avoid double-counting. Consequently, any packets sent directly (not via a policer) to a local SAP post-policer queue are not counted in the **sap-stats** output. The output of **show service id** *service-id* **sap** *sap-id* **stats** always counts these packets in both the related policer and queue. If it is required to count packets sent directly to the local SAP post-policer queue in the **sap-stats** output, the packets could be sent into a policer with the rate set to maximum and then into the local SAP queue.
- When redirecting traffic in a SAP egress QoS policy to a policer in an **ip-criteria** or **ipv6-criteria** statement, without **use-fc-mapped-queue** being configured in the criteria **action** statement, the redirected traffic of a subscriber with an *inter-dest-id* matching a configured **host-match** statement (under an egress queue group or Vport) will exit via the port's default *policer-output-queues* queue group instead of the egress queue group associated with the **host-match** statement. [236293]
- The following features are not supported on an IOM4-e-HS:

  - H-QoS Virtual Scheduling, (IOM4-e-HS equivalents are available) which includes:
    - card-based virtual scheduler adjustments
    - egress queue and policer parenting, and the associated overrides, to schedulers
    - parenting to a port-scheduler
    - egress non-IOM4-e-HS aggregate rate limiting
    - advanced configuration policies
    - limit unused bandwidth
  - Queue dynamic MBS, CBS, drop tail and burst limit commands and their overrides (an IOM4-e-HS equivalent is available for the burst limit)
  - Queue CIR and CIR adaptation rules and their overrides
  - Queue aggregate rate frame based accounting
  - Queue average frame overhead
  - Egress network queue MBS setting (an IOM4-e-HS equivalent is available)
  - WRED per queue (an IOM4-e-HS equivalent is available)
    - pool-per-queue mode
    - native mode
  - The highplus slope in a slope policy
  - Ingress shared queuing and multipoint-shared queuing
  - All HS-MDA related commands, including **exp-secondary-shaper**, but excluding **pw-sap-secondary-shaper**
  - Egress regular pools, HS-MDA pools, and named pools (IOM4-e-HS equivalents are available)
  - Ingress buffer reallocation
  - Stable pool sizing (an IOM4-e-HS equivalent is available)
  - All Vport related commands
- Changing the CBS of queues on FP4-based line cards might result in a very small loss of packets.

11.47 Filter Policies

- QoS and IP filter matches on IP frames are limited to Ethernet Type II IP frames. In particular, Ethernet SNAP IP frames will not be matched with IP match criteria. [15692]

- IP filters with a **default-action drop** will not drop non-IP packets (such as ARP and IS-IS). [40976]
- MAC filtering does not match on IPv6-enabled IES interfaces. [44897]
- The HTTP-redirect action is allowed in MAC-filter policy configurations, but the action is not supported for MAC-filter policies. [140058]
- Configuration rollback may fail when rolling back changes on filters with entries overwriting embedded-filter entries if the filter configuration at any stage of the rollback exceeds the supported filter configuration limits. This can only happen when the embedded filter entry and the embedding filter entry require different hardware resources. [162867]
- A CPM filter policy does not support an **action-queue** for VRRP protocol match but this configuration is not blocked in CLI. [164497]
- For VPRN services that use GRE tunnels as transport, applying an egress **ip-filter** on the network interface of the originating node will match fields of the inner IP header and not the outer GRE IP header. [189799]
- The existing filter policy functionality does not provide notification when a PBR/PBF redirect changes either as result of PBR target going down or being deleted, or as a result of PBR target reprogrammed for a redirect policy. [198852]
- Adding or removing prefixes within an existing **match-list ip-prefix-list**/**ipv6-prefix-list** context, referenced by CPM and/or IOM filter **src-ip** and **dst-ip** match criteria, can result in a small amount of packet loss. The same is applicable when adding or removing port(s) or port range within an existing **match-list port-list** context, referenced by CPM and/or IOM filter **src-port/dst-port** and **port** match criteria. [257606]

11.48 PBR/TCS

- If a Transparent Cache Switching (TCS) redirect-policy destination does not have a test clause defined, the operational state is reported as “Up”. [21227]
- An IP address must be assigned to the system interface and the interface must be operationally up in order for Web portal or HTTP redirect to operate. [46305]
- The Nuage Service Chaining for IES/VPRN using IPv4 filter ESI PBR for EVPN feature has the following known limitations.
    - Only unicast traffic is subject to PBR; other traffic matching a Layer-3 ESI PBR entry will be subject to action forward.
    - The egress EVPN interface must be in a VPRN service (same or different routing instance).

  - The Service Function appliance must be in the local IP subnet reachable via the specified EVPN egress interface.
- The PBR feature ESM downstream traffic steering using egress IPv4 ACLs with PBR action has the following limitations.
  - Only unicast traffic is subject to L3 PBR; any non-unicast traffic matching a Layer-3 entry will be subject to action forward. The same rule applies to traffic matching a filter entry with an egress PBR action if the filter is deployed in the ingress direction.
  - Local-to-local subscriber/host traffic when both subscribers are subject to VAS scenario is not supported in production networks.

11.49 Services General

- The CLI does not display an error when the user attempts to apply a filter log and a mirror-source to a given SAP at the same time. A filter log and mirror-source cannot be applied simultaneously to the same SAP. [22330]
- When the standby spoke-SDP of an endpoint becomes active due to a revert-time expiration or a forced switchover, the Multi-Tenant-Unit (MTU) SAP may forward duplicated packets (only of broadcast/multicast/unlearned unicast types) coming from the redundant spoke-SDPs for a few milliseconds. For broadcast TV distribution and similar applications where the duplicated packets may have a side-effect, Nokia recommends that the redundant spoke-SDPs be operated in non-revertive mode. [67252]
- If a configuration is saved (**admin save**) after enabling the MC-ring status by **no shutdown** and the related configurations such as SRRP, BFD and IBCP are modified and cause a “CONFIG_ERR” in MC-ring afterwards, the saved configuration may have reloading issues. [78245]
- If an MC-ring breaks, slow RNCV is not performed and fast RNCV stops the moment one of the peer detects the ring node. The ring node that detects the peer first receives the connected status. [78246]
- When the **ce-address-discovery** option is enabled on an Ipipe VLL service and the Ethernet SAP comes back up from an operationally down state due to link failure, the PE node will forward IP multicast/broadcast packets over the Ethernet SAP but drops IP unicast packets until an ARP message is received from the CE router. This is in accordance to *draft-ietf-l2vpn-arp-mediation*. When the Ethernet VLAN SAP is switched through an Ethernet switch or NTE device that does not implement Ethernet OAM fault propagation, the CE node may not be aware of the link failure and will not generate an ARP message to update the PE ARP cache until the time when the ARP cache in the CE times out. The only workaround is to set the ARP cache timeout to a lower value on the Ethernet CE router. [78805]

- A Multi-Site Scheduler (MSS) must either have a single (card-level) scheduler hierarchy instantiated, or have a scheduler-hierarchy instantiated per member port for multi-member logical ports such as LAG and APS, but not both. When an APS SAP is added to an MSS, a site_instance is created for each APS group member port, and a scheduler hierarchy is instantiated per site instance. If a regular (physical port) SAP was also to be added to the same MSS, then a card-level scheduler hierarchy would be created. The per site-instance scheduler hierarchies and the card-level scheduler hierarchy within the MSS are disconnected and therefore would not provide a meaningful H-QoS function. [81279]
- A GRE SDP is not supported over an RSVP shortcut. The GRE SDP will go down if the destination is reachable via an RSVP shortcut route. [91257]
- For Distributed CPU Protection, the rate limiting is per-protocol per-SAP (or per network interface). It does not support rate limiting per individual subscribers within a single SAP. This limitation also applies to capture SAPs. All control traffic for subscribers that have not yet established an MSAP is treated as a single aggregate (per protocol). Configuration is via CLI and SNMP; there is no RADIUS support.
- Configuration of IPv6 is not supported on Ipipe spoke-SDP terminations in an IES or VPRN service context. [128543]
- The following features are not supported on EVPN-enabled Routed-VPLS interfaces in VPRN services: IS-IS, RIP, OSPF, and authentication-policy. [168271]
- An R-VPLS interface binding to a VPLS service will make the R-VPLS interface operationally down if the R-VPLS interface MAC-address matches a static-MAC or OAM-MAC configured in the associated VPLS service. In this scenario, to restore the R-VPLS interface to be operationally up, either one of the following actions need to be taken:
  - Change the R-VPLS interface MAC-address
  - Remove the conflicted static- or OAM-MAC address and then unbind and re-bind the R-VPLS interface configuration. [170516]
- For R-VPLS, configuring **service-mtu** to a value lower than 142 will result in packets exceeding the configured **service-mtu** value being dropped with no IP fragmentation. [180872]
- Support of XMPP on a DC PE in VPLS/VPRN requires the user to use all lowercase letters while configuring the username field with **configure system xmpp server** *xmpp-server-name* **create username** *user-name* **password** *password* **domain-name** *domain-name*. The CLI/SNMP does not reject configuring any uppercase letters, but only lowercase letters are functionally supported. This is due to ejabberd (Erlang Jabber Daemon) interoperability issues and how ejabberd interprets uppercase user names. [190076]

- EVPN IP routes will not be added to the RTM if the VPRN service is operationally down, except if it is down because of a missing route-distinguisher configuration. [192237]
- VCCV BFD is not supported on MPLS-TP PWs (that is, where **pw-path-id** is configured).
- BFD sessions, where the BFD Template specifies type **cpm-np**, are not supported by VCCV BFD.
- The following limitations apply for Pseudo-Wire SAPs (PW-SAPs):
  - PW-SAPs require IOM3-XP/-B/-C and are supported with the HS-MDAv2
  - PW-SAPs are only supported on Epipe VLL services, as well as on interfaces and group interfaces in an IES or VPRN service.
  - Only Ethernet PWs are supported
  - Ethernet CFM is not supported on the Ethernet PW or PW-SAP
  - No support for mixed SDP types
  - No support for PW control word
  - No support for hash-labels
- The XMPP support on DC PE for the VPLS/VPRN (Fully-Dynamic model) feature is not supported in combination with the RADIUS-triggered dynamic data services feature in the same system. The two features are mutually exclusive.
- For XMPP support on a DC PE for the VPLS/VPRN Fully-Dynamic model, when the VSD creates a configuration in the system, rollbacks could fail in those situations where policies are created by CLI/SNMP but the association to services is provisioned by the VSD.
- Protocol classification and identification of underlying functions are not supported at either ingress or egress for frames received at ingress with more than two VLAN tags.
- The configuration of Epipe services is not supported from VSD through the Fully-Dynamic integration model, although Epipe commands are shown in the **tools dump service vsd-services** command-list. [217287]
- The router policy statements "_ES_EvpnEthSegRtExp" and "_ES_EvpnEthSegRtImp" are auto-created by the system for EVPN multihoming functions. It is advised not to use these policy statements in any configuration contexts, as they are reserved by the system. [218217]
- Assuming **force-vlan-vc-forwarding** is configured in a PW-template being used by BGP-AD, when **provider-tunnel** is enabled and its owner is **bgp-ad**, the root node does not preserve the ingress tag. [218480]
- Black-hole MAC addresses are not supported in B-VPLS services.
- The following constraints must be considered when configuring **connection-profile-vlan** SAPs:
  - Not supported in the following type of services:

        - Etree
        - M-VPLS
        - B-VPLS
        - R-VPLS
        - I-VPLS
        - PBB-Epipe
    - The following features are not supported in combination with **connection-profile-vlan** services:
        - **proxy-arp** and **proxy-nd**
        - Capture SAPs
        - **eth-tunnel** SAPs
        - **eth-ring** – Connection-Profile (CP) SAPs can be used as th-ring data SAPs but control G.8032 traffic is not supported in CP SAPs.
        - **vlan-translation**
        - xSTP – CP SAPs can be managed by an M-VPLS, but services with CP SAPs do not support xSTP.
        - L2PT
        - BPDU translation
        - Subscriber management features
        - IGMP/MLD/PIM (v4 or v6)
        - **vlan-vc-tag** under an SDP-binding sharing service with a CP SAP.
    - In Release 14.0.R1, ETH-CFM configuration is restricted on CP SAPs with the exception of the ETH-CFM **vmep-filter** option. The configuration of vMEP-filters on CP SAPs is highly recommended in services where vMEPs are configured so that all untagged ETH-CFM traffic cannot be leaked out of the CP SAPs.
- **bgp-evpn mpls force-vlan-vc-forwarding** is not supported on R-VPLS services. In addition, a configuration file containing **force-vlan-vc-forwarding** and **provider-tunnel** leaf-only configuration (that is, **no root-and-leaf**) in an EVPN R-VPLS service will fail to execute. [228492]
- The following features are not supported for PW-ports bound to physical ports:
    - PW with GRE transport
    - VPLS service
    - NULL PW-port encapsulation
    - **vc-type vlan** together with QinQ PW-Port
    - VCCV-BFD
    - PW control word

  - Hash labels
  - Entropy labels
- The following features are not supported on FPE-based PW-ports:
  - PW-port that is associated with an FPE cannot be part of Multi-Service-Site (MSS)
  - VPLS service
  - NULL PW-Port encapsulation
  - VCCV-BFD
  - PW control word
  - MC-LAG
- In case of GREv6 over IPsec, the operator needs to use **dest-ip** configuration under IPsec tunnel to resolve the GRE peer address; using static-route with IPsec tunnel next-hop is not currently supported. [234668]
- If a BGP export policy is used to change the local preference of BGP-VPLS and BGP multi-homing updates on a system advertising these updates to an EBGP peer, the VPLS preference in the Layer-2 info extended community in these updates will not be set to the modified local preference value. This could cause a system in a remote AS to receive the same update with different VPLS preference values if the updates are received over different EBGP peering sessions. [256401]
- The following features are not supported on an IOM4-e-HS:
  - Customer multi-service sites
  - G.8031 protected Ethernet tunnels
  - PBB egress B-SAP per-ISID shaping
- MACsec has the following limitations:
  - Up to 36 ports with MACsec are supported per chassis, although there is no enforced hard limit.
  - On a LAN configuration (multiple MKA peers on the same interface), MACsec is not supported on production networks.
  - For an XPN cipher-suite (extended packet number), MACsec is not supported on production networks. XPN is designed to be used on high-speed interfaces (100Gbps) as a more efficient way of counting MACsec packets so Security Association Keys (SAKs) do not have to be updated as often.

11.50 EVPN Multihoming

- SAPs/SDP-bindings belonging to a given **ethernet-segment** but configured on non-BGP-EVPN-MPLS-enabled VPLS or Epipe services will be kept operationally down with the StandByForMHProtocol flag.
- Null Ethernet ports are not supported on virtual Ethernet Segments (vES).
- **connection-profile-vlan** SAPs cannot be associated with a vES and cannot be configured on ports where vESs are defined. They may, however, be configured on different ports on the same service.
- The association of a PW-port to an ES or vES is not supported.
- Ports where **eth-ring** SAPs are defined cannot be added to Ethernet Segments or virtual Ethernet Segments. [264461]
- SAPs defined in an MC-rings cannot be added to Ethernet Segments or virtual Ethernet Segments. [264461]

11.51 PBB-EVPN

- When **bgp-evpn mpls** is enabled in a B-VPLS service, an I-VPLS service linked to that B-VPLS cannot be an R-VPLS (the **allow-ip-int-bind** command is not supported).
- The ISID value of 0 is not allowed for PBB-EVPN services (I-VPLS and Epipes).
- The following features/commands are not supported in an I-VPLS when **bgp-evpn mpls** is configured in the B-VPLS service:
  - **mac-protect** and **auto-learn-mac-protect**
  - **end-point** and **attributes**
  - **eth-tunnel**s
  - sharing of ports or SDPs between a B-VPLS service enabled with **bgp-evpn mpls** and its associated I-VPLS/Epipe services is not allowed.
- EVPN all-active multi-homing is not supported within a B-VPLS configured for EVPN-MPLS when PIM snooping is enabled in an associated I-VPLS. [251610]
- For PBB-EVPN E-Tree:
  - BGP-MH sites are not supported on I-VPLS E-Tree services
  - EVPN all-active multi-homing is not supported on I-VPLS leaf Attachment Circuit (AC) SAPs

11.52 PBB-EVPN Multihoming

- Ethernet Segments (ESs) can be associated with B-VPLS SAPs/SDP-bindings and I-VPLS/Epipe SAPs/SDP-bindings; however, the same ESs cannot be associated with B-VPLS and I-VPLS/Epipe SAP/SDP-bindings at the same time.
- When PBB-Epipes are used with PBB-EVPN multihoming, the following restrictions apply:
    - PBB-Epipe spoke-SDPs are not supported on ESs.
    - For non-local-switching PBB-Epipes (there is a single SAP per Epipe) only all-active multihoming is supported.
    - For local-switching-enabled PBB-Epipes (two SAPs are defined within the PBB-Epipe instance):
        - only single-active multihoming is supported
        - only when the two ends of the PBB-Epipe are defined in two systems (and not three or more)

11.53 QinQ Default SAPs

- The following constraints must be considered when configuring *.null and *.* QinQ SAPs:
    - only supported in Ethernet ports or LAG
    - only supported on Epipe, PBB-Epipe, VPLS and I-VPLS services. They are not supported on VPRN, IES, R-VPLS or B-VPLS services.
    - capture SAPs with encapsulation :*.* cannot co-exist with a default :*.* SAP on the same port
    - inverse-capture SAPs (*.x) are mutually-exclusive with :*.null SAPs
    - no support for:
        - PW-SAPs
        - **eth-tunnel** or **eth-ring** SAPs
        - **vlan-translation copy-outer**
        - E-tree **root-leaf-tag** SAPs
        - subscriber-management features
        - BPDU-translation
        - IGMP-snooping
        - MLD-snooping

- ETH-CFM primary-VLAN

11.54 Subscriber Management

- Dynamic subscribers learned (via DHCP) while **sub-sla-mgmt** is shut down will continue to use the SAP-level ingress and egress filter rules. Once the subscriber is relearned (renewed), the subscriber profile filters will then be used. This does not apply to static subscribers. [47167]
- Since the SR routing model is based on a broadcast Ethernet network, the IP addresses of the subnet (for example, x.y.0.0/16 or x.y.z.0/24) and the subnet broadcast address (for example, x.y.255.255/16 or x.y.z.255/24) should not be used as IP addresses for both IPoE (DHCP/static/ARP) subscribers. PPPoE hosts can use these addresses starting from Release 9.0.R3 with the support for PPPoE unnumbered interfaces. [78233]
- When a CoA request is sent for changing the subscriber-ID of a subscriber host in a dual-stack PPPoE session, both the IPv4 and IPv6 hosts will have their information changed. This may temporarily increase the subscriber count on the SAP, which should be reflected in the **multi-sub-sap** limit. [90556]
- When a RADIUS CoA message triggers a change of subscriber-profile and/or username together with a change of SLA-profile, a RADIUS Accounting-Stop message or RADIUS Accounting-interim-update message (reason sla-stop) is generated for the subscriber. These accounting messages do not include the old subscriber-profile name and/or old username, but only those from the CoA message. [94758, 256628]
- In a network where DHCP relay is dual-homed, a VPLS SAP with DHCP-snooping enabled will receive two identical DHCP reply messages from the DHCP server. When RADIUS authentication is enabled on the VPLS SAP and the DHCP server did not echo the Option 82 information, RADIUS authentication will be executed again for DHCP reply messages. For dhcpACK messages, if the SR OS still has an outstanding RADIUS transaction from the first dhcpACK when receiving the second dhcpACK, the latter one will be dropped and a dhcpRelease message will be incorrectly generated towards the DHCP server. When RADIUS authentication is successful for the first dhcpACK, the client will still receive the dhcpACK and starts using the IP address. [101767]
- Direct replication over subscriber hosts in the subscriber management context has been extended to support replication to two new modes, but have the following limitations:

  - Per SAP replication — in this mode, only a single copy of a multicast stream per SAP is transmitted regardless of the subscriber management deployment model (subscriber per SAP, service per SAP or a single SAP per all subscribers). For example, if multiple hosts on a SAP are subscribed to the same multicast group, only a single copy of multicast stream will be sent towards the access network. In this model, multicast traffic is flowing outside of the subscriber queues. IGMP states are maintained per host and SAP.
  - Multicast traffic can be redirected to a different interface from the interface on which IGMP join has arrived. Redirection is supported within a VRF, within the GRT and between VRFs. However, redirection between the GRT and a VRF (and vice versa) is not supported. Multicast redirection is a new feature and should not be confused with host tracking although the functionality of the two are very similar. Host tracking is still supported. For a given subscriber, the usage of IGMP and host tracking is exclusive; they cannot both be active on the same subscriber.
- When a subscriber host makes use of policers feeding into queues, the queuing statistics require the reconciliation of the policer and queue statistics. Therefore, Nokia recommends waiting at least 10 seconds after traffic has stopped before issuing a **clear statistics** command. [115390]
- The following ESM Multi-Chassis Sync (MCS) client applications are not blocked in CLI but should not be enabled in MCS on hybrid ports in production networks: **igmp**, **igmp-snooping** and **mld-snooping**. [123469]
- When using **host-lockout** on managed SAP's using one VLAN for all PPP sessions, some sessions can become locked-out during the initial setup in case of high setup rates [126348]
- In case a QinQ capture SAP has a port inner Ethernet type value configured different from the default value “0x8100”, and **authentication-policy** uses **pap-chap** as **pppoe-access-method**, the PPPoE PADO message is incorrectly sent out from the MSAP with the default inner Ethernet-type 0x8100. This is not an issue in case the capture SAP is dot1q-tagged or the **authentication-policy** uses a different **pppoe-access-method**. [137800]
- The following restrictions for DHCPv4 over PPPoE apply:
  - The DHCPv4 client must be connected via a CPE that acts as a DHCP relay.
  - The DHCPv4 client subnet must be known as a managed route attached to the subscriber PPPoE host (next-hop of managed route is the PPPoE host)
  - The DHCP Relay Agent IP address (giaddr field) inserted by the CPE DHCP relay must be part of the managed route subnet (not the subscriber PPPoE host’s IP address)

  - Downstream DHCPv4 over PPPoE frames will be sent through the egress SLA instance queues of the PPPoE subscriber; hence, they are part of the subscriber QoS scheduling context. [137283, 138115, 138890]
  - The DHCP server is not local on the node where the PPPoE/LNS session is terminated. [138242, 138972]
- An SR OS-based DHCPv6 server can only be used in combination with an SR OS-based DHCPv6 relay on a group interface with Enhanced Subscriber Management (ESM) enabled or with an SR OS-based DHCPv6 relay on a regular service interface. Using an SR OS-based DHCPv6 server as a standalone server with a non-SR OS-based DHCPv6 relay is not supported. [149028]
- The following restrictions apply for the Wholesale/Retail routed-CO model:
  - An up-front Layer-3 DHCPv4 or DHCPv6 relay agent in combination with Wholesale/Retail configuration is not supported. [72138]
  - Leaking of a subscriber prefix from a retailer VPRN into a different local VPRN or leaking static, managed or BGP routes that have a subscriber prefix as next-hop is not supported. [134840, 140643]
  - No support for static IPv4 hosts on unnumbered retail subscriber interfaces [150733]
  - Synchronization of subscriber IGMP/MLD states between redundant BNG nodes protected via the same MC-LAG/SRRP protection mechanism and part of a Wholesale/Retail setup is currently not supported. The IGMP/MLD state will be synchronized to the standby node but will fail installation with the reason "IGMP/MLD interface not found". [155540]
  - ESM multicast enables ESM group interfaces to process each host's IGMP and/or MLD messages; and hence, enabling IPv4 or IPv6 multicast delivery to individual ESM host. ESM multicast is supported only if both the Wholesale and Retail are VPRN services. ESM multicast is not supported if the Retail is an IES instance. [179941]
  - Overlapping addresses in retail services (**private-retail-subnets**) are supported for PPPoEv4, PPPoEv6 and IPoEv6. They are not supported for IPoEv4. [191027]
  - No multi-chassis redundancy support in combination with overlapping addresses in retail services (**private-retail-subnets**)
  - IES as a retail service is not supported for IPoEv4 hosts
  - No support for PPPoA and PPPoEoA sessions
  - Unique IPv4 subnet per subscriber for IPoE (**virtual-subnet**) is not supported in a retail service.
  - Web Portal Protocol (WPP) is not supported on a retail subscriber interface
- L2TP tunnels over LDP shortcuts are not supported. [154574]

- The initial DHCP message of an internal DHCPv4 client for PPPoE requests a **lease-time** of one hour. However, the next DHCP renew or rebind will use the last granted **lease-time** from the DHCP server. If the granted **lease-time** was equal to the Maximum Client Lead Time (MCLT) because of a **local-dhcp-server** used in **failover** mode, Nokia recommends enforcing at least the default **lease-time** of one hour by configuring the pool **min-lease-time**. [157485]
- Although “FRAMED INTERFACE ID” is configured below the RADIUS Accounting policy, the parameter can be missing in the Accounting-Stop message for certain termination root causes such as “User Request(1)” and “Admin Reset(6)”. This is not an issue for termination root cause “Lost Carrier(2)”. [164568]
- ECMP load-balancing to identical RADIUS Framed-Routes/Framed-IPv6-Routes with different next-hop is not supported in the following Wholesale/Retail scenario:
  - A combination of ECMP Framed-Routes/Framed-IPv6-Routes belonging to hosts on a subscriber interface with **private-retail-subnets** enabled and hosts on a subscriber interface without **private-retail-subnets** enabled.

  In this scenario, a part of the ECMP load-balanced traffic is dropped. [167136]
- Setting up a Diameter peer TCP connection via VPRN is only supported with the default TCP port 3868. [186325]
- A setup with an up-front DHCP relay server and having **lease-populate l2-header** enabled on the second relay that is part of the same routing instance as the **local-dhcp-server**, is not supported. The workaround is to have the **local-dhcp-server** external to or in a different routing instance than the second relay. [192649]
- Persistency file sizes larger than 2GB are not supported. When a persistency file reaches the 2GB file size limit, an event is raised and persistency will stop saving data to the compact flash. An operator intervention is required to re-initialize the persistency file using the following CLI commands: **config system persistence** *client-application* **no location** followed by **config system persistence** *client-application* **location** *cflash-id*. [199023]
- IPoE IPv6 hosts that share a /64 prefix (**ipoe-bridged-mode**) with separate **sla-profile** instances are not supported. Egress traffic for these hosts will share a single (arbitrary) set of **sla-profile** instance queues/policers. [199934]
- A configuration rollback can fail when a static IPv6 host is configured on group interface SAPs [200715]
- The oversubscribed multi-chassis redundancy model in ESM has the following limitations:
  - Central standby node must use SF/CPM4 or higher (other protected nodes can continue to use SF/CPM3).
  - All nodes in the OMCR cluster (central standby and the protected nodes) must run Release 12.0.R1 or higher.

  - While the node is in the central standby mode of operation, the configuration of 1:1 (active-active) peering session on the same node is blocked. In other words, the central standby mode of operation becomes the only mode of operation on that node.

    However, non-central standby nodes can have a peering connection with a central standby backup node (OMCR mode of operation) and at the same time another peering connection with another active BNG node in the 1:1 model.
  - Only DHCPv4/v6 subscribers in the Routed Central Office (RCO) model are supported.
  - Synchronization of the following MCS clients is not supported:
    - Host tracking
    - MC-ring
    - Layer-2 subscriber hosts
    - Layer-3 IGMP/MLD
    - Layer-2 IGMP/MLD
    - DHCP Server
    - PPPoE Clients
    - MC-LAG
    - MC-IPsec
    - MC-endpoint
  - The failover trigger is based on SRRP only (no MC-LAG support).
  - Pre-emption of already instantiated subscriber hosts in the central standby node by another subscriber hosts is not allowed.
  - Persistency in multi-chassis environment must be disabled since redundant nodes are protecting each other and they maintain up-to-date lease states.
- An IPv6 subscriber can be mirrored/LI'd using the subscriber ID as the mirror/LI source criteria, but a specific IPv6 host cannot be a source criteria (only the subscriber which will include all IPv6 hosts associated with that subscriber ID).
- The maximum number of hosts within the subscriber or the SLA-profile instance that can be affected by a single CoA is 32.
- IPoE hosts with separate SLA-profile instances and duplicate MAC addresses on a single SAP with **nh-mac** antispoofing are not supported. Ingress traffic for these hosts will share a single (first created) set of SLA-profile instance queues. This restriction has been in place since Release 6.0.
- BGP peering between CPE and BNG via a managed route is not supported.
- An SR OS-based DHCPv6 relay on a regular interface cannot be used in combination with **antispoof ip/ip-mac/mac** on the SAP.

- An SR OS-based DHCPv6 relay on a regular service interface cannot be used in combination with an authentication policy on that interface.
- Diameter NASREQ authentication is not supported
  - for L2TP LAC hosts nor L2TP LNS hosts
  - on group interfaces of type **lns** or **wlangw**
- The following restrictions apply for IPoE sessions:
  - ARP hosts are not supported in an IPoE session and cannot be instantiated on a group interface with IPoE sessions enabled.
  - A local user databaase host identification based on option 60 is ignored when authenticating an IPoE session.
  - RADIUS authentication of an IPoE session fails when the **user-name-format** is configured to **mac-giaddr** or **ppp-user-name**.
  - The alc.dtc.setESM() API in the DHCP Transaction Cache (DTC) Python module cannot be used in combination with IPoE sessions.
  - The DHCP Python module (alc.dhcp) used to derive subscriber host attributes from a DHCPv4 ACK message is not supported in combination with IPoE sessions.
  - WPP is not supported in combination with IPoE session.
  - The creation of an IPv4 host using the Alc-Create-Host attribute in a RADIUS CoA message is not supported on a group interface with IPoE session enabled.
  - A RADIUS CoA message containing an Alc-Force-Nak or Alc-Force-Renew attribute is not supported for IPoE sessions.
- The following restrictions apply for Layer-3/IP accounting:
  - Layer-3/IP accounting is not supported in combination with last-mile-aware shaping on HS-MDAv2 MDAs
  - Layer-3/IP accounting is not supported in combination with ESMoPW on HS-MDAv2 MDAs
  - Layer-3/IP accounting is not supported in combination with MLPPP
  - Layer-3/IP accounting in combination with ESMoPW and last-mile-aware shaping may be inaccurate if the MPLS encapsulation overhead changes during the lifetime of a subscriber.
  - Layer-3/IP accounting is restricted to a single encapsulation per SLA-profile instance (queue instance). The first host associated with the SLA-profile instance (queue instance) determines the allowed encapsulation. Conflicting encapsulations are:
    - PPPoE and IPoE on regular Ethernet SAPs
    - PPPoE and IPoE on PW-SAPs
    - PPPoA and PPPoEoA on ATM ports

  - PPPoE keepalive packets do not contain IP payload and introduce an error in Layer-3/IP accounting when enabled in combination with L2TP LAC. A workaround is to isolate the keepalives in a separate queue/policer.
  - Padding of frames smaller than the Ethernet minimum frame size (64 bytes) may introduce an inaccuracy in Layer-3/IP accounting.
  - With ATM in the last mile, last-mile-aware shaping may introduce an inaccuracy in Layer-3/IP accounting.
  - Packet-Byte-Offset (PBO) changes during the lifetime of a subscriber introduces an inaccuracy in Layer-3/IP accounting.
- On HS-MDAv2, there is no per-egress queue granularity to count IPv4- and IPv6- forwarded/dropped subscriber traffic separately. When stat-mode v4-v6 is configured on an egress HS-MDAv2 queue, it is applied to all egress queue-group queues for that subscriber.
- **mac-sid-ip** anti-spoofing for PPPoE on the group interface cannot be used in combination with L2TP LAC.
- ESM is supported on the 4-port 100GE CXP, 4-port 100GE CFP4, and 40-port 10GE SFP+ IMMs with the following restrictions:
  - static SAPs (non MSAP) are only supported with policers on ingress
  - MSAPs are now supported, but with the following limitations:
    - **profiled-traffic-only** is mandatory for MSAPs on this type of IMM
    - **msap-policy** needs to define ingress service-queueing because ingress shared-queueing is not supported
    - no support for multiple subscribers per MSAP due to the mandatory **profiled-traffic-only** setting which is a **single-sub-parameter**
    - no support for per-SAP multicast replication into the MSAP because of the **profiled-traffic-only** setting
- Diameter multi-chassis redundancy is not supported for OMCR (Oversubscribed Multi-Chassis Redundancy). Diameter applications (Gx, Gy, NASREQ) in general are not supported in combination with OMCR. Gx for Usage-Monitoring and AA is currently not supported in multi-chassis configurations.
- Stateful MC-LAC redundancy does not protect tunnels against a node failure for **failover recovery-method mcs**, introduced in Release 13.0.R1. Stateful MC-LAC redundancy does protect tunnels against a node failure for **failover recovery-method recovery-tunnel**.
- In case the same L2TP tunnel client endpoint is shared by LAC sessions under multiple group interfaces, then all SRRP instances need to share fate using an **oper-group**: all SRRP instances in the group will switch together to the redundant LAC when an error is detected.

- When LAC sessions under a group interface are spread over multiple LAC tunnels with different L2TP tunnel client endpoints, all interfaces used for LAC tunnel client endpoint addresses need to track the same SRRP instance for fate sharing.
- ESM **host-lockout** is not supported for LNS.
- When using Python-policy cache persistency on the 7750 SR-a4/a8, a persistency-downgrade to Release 12.0.R9 or 13.0.R1 is not supported. [201175]
- When multiple identical framed routes are received for a single subscriber host or IPoE/PPP session, only the first framed route will be accepted while all subsequent identical framed routes are silently ignored. Framed routes are considered identical when prefix and prefix length are the same, irrespective of the specified metrics. This applies to both IPv4 and IPv6 framed routes. [205607]
- For 7750 SR-7/12/12e, 7450 ESS-7/12 and 7450 ESS-7/12 mixed mode chassis types, the unnumbered DHCPv6 IA-NA subscriber hosts are limited to 128k per system, or to 64k per system in case the unnumbered DHCPv6 subscriber hosts are terminated on a retail subscriber interface (Wholesale/Retail). This limit is not enforced by the system. Unnumbered DHCPv6 IA-NA subscriber hosts are those that have a prefix that falls outside the provisioned subscriber WAN-host prefixes on the subscriber interface. Support for unnumbered subscriber hosts must be explicitly enabled per subscriber interface with the **allow-unmatching-prefixes** CLI command for IPv6. [206968]
- When DHCPv6 IA-PD is modeled as a managed route pointing to an IPv4 subscriber host as next-hop (**pd-managed-route next-hop ipv4**), the following restrictions apply.
  - There are no ingress or egress IPv6 filters installed for traffic from/to the PD prefix.
  - There are no ingress or egress QoS IPv6 criteria installed for traffic from/to the PD prefix.
  - Multicast replication to the PD prefix is not supported. [209165]
- For GRT lookup and Routed-CO VPRN, exporting **sub-mgmt** and **managed** routes from a VPRN service to GRT leak when **srrp-enabled-routing** is configured in the subscriber management group interfaces may result in routing instabilities and black-holing during CPM/CFM activity switchover events. The exporting of these route types in a dual-homed scenario should be avoided. [220779]
- The following restrictions apply for RADIUS Subscriber Services.
  - Subscriber services are not synchronized in Multi-Chassis Synchronization (MCS). They must be re-applied after failover in a multi-chassis redundant deployment.

  - **rate-limit** (PIR/CIR) and account actions are not supported in PCC-rule subscriber services on L2TP LNS sessions.
  - An egress rate-limit (PIR/CIR) action is not supported in PCC-rule subscriber services on HS-MDAv2. Egress dynamic policers on HS-MDAv2 are always installed with PIR=max and CIR=0.
  - On HS-MDAv2 only a single SLA-profile instance can be active for a subscriber when a PCC-rule subscriber service contains egress QoS actions. A PCC-rule subscriber service with egress QoS actions must be removed before the SLA-profile of an HS-MDAv2 subscriber can be changed.
  - There is no support for hierarchical policing on HS-MDAv2 egress dynamic policers that are instantiated by PCC rule-based subscriber services.
  - PCC-rule subscriber services are not stored in the **subscriber-mgmt** persistency file.
  - RADIUS PCC-rule subscriber services and Diameter Gx-provisioned PCC rules cannot be provisioned simultaneously for the same PPPoE or IPoE session.
- The following restrictions apply for PCC rules (both initiated from Gx and RADIUS subscriber services).
  - PCC rules are not supported on L2-Aware NAT hosts.
  - PCC-rules use CAM resources in filter and QoS policies. Careful planning and a high degree of policy and rule sharing is required for a scalable deployment.
  - PCC-rules are not supported on L2TP LAC sessions, PPPoEoA, PPPoA, MLPPP, non-sub-traffic hosts and static-hosts. These host types should not be part of an SLA-profile instance where other subscriber hosts or sessions with PCC rules are active.
  - When an SLA-profile instance contains multiple subscriber hosts, it is mandatory that all hosts have the same PCC rules applied.
- The following restrictions apply for data-triggered subscriber management:
  - L2-Aware NAT is only supported on vRGW-enabled interfaces.
  - Subscriber hosts must be on Ethernet ports on IOM3-XP/-B/-C/IMM or higher.
  - Data-triggered host setup and promotion is only supported when AAA/ LUDB returns all of the IP information. For promotion to DHCP relay, the Alc-Force-DHCP-Relay VSA must be included.
  - Data-triggered promotion is not supported with **anti-spoof nh-mac**.

  - Data-triggered host setup and promotion with IPoE **session-key sap | mac | cid** and **user-ident mac-interface-id** is only supported when AAA returns the Agent-Circuit-Id VSA and the **circuit-id-from-auth** flag is set in the IPoE session policy (IPoE session merging).
  - Unnumbered data-triggered host setup is only supported when the SR OS node receives Framed-IP-Netmask from AAA or LUDB.
  - Alc-Force-DHCP-Relay VSA is not supported with Gx, NASREQ and LUDB. It must be returned via RADIUS.
  - Inter-SAP mobility and all SHCV triggers are only supported for the same host-type. Data-triggers cannot trigger SHCV checks for a DHCP lease-state and DHCP cannot trigger SHCV checks for a data-triggered host.
  - Diameter Gx/Gy CCR-Ts are not sent after SRRP switchover on a IPoE session stateless redundancy group interface.
  - Data-trigger promotion is not supported for Wholesale/Retail.
  - Data-trigger promotion is not supported on **unnumbered** / **allow-unmatching-subnets** / **allow-unmatching-prefixes** subscriber interfaces.
- Diameter Gx is not supported on static hosts, L2TP LAC sessions, and MLPPP sessions on LNS.
- Diameter Credit Control (Gy) is not supported on L2TP LAC hosts.
- The credit control category map specified in a CCA-I Charging-Rule-Base-Name AVP is ignored when Diameter Gy Extended Failure Handling (EFH) has been active. [233571]
- For MCS/SRRP with BNGs running different versions, MC-ring is not supported between redundant BNGs running different versions.
- DNAT-only is not supported in a dual-homing configuration.
- Downstream L2TP LAC traffic redirection (network router to VAS) over R-VPLS interfaces does not work when LAGs are configured on the network interface and only a single R-VPLS next-hop IP address is supported (SR OS only redirects to one VAS). [249132]
- All managed routes (including **pd-managed-route**) of IPv4 data-triggered ESM hosts will be withdrawn and re-advertised upon promotion to a DHCP ESM host. [250763]
- The sum of the number of native L2TP LAC tunnels and the number of L2TP LAC VAS tunnels (for steered sessions) may not exceed 16K-1 in case of unique tunnel peers (unique LNS server per tunnel). As a result, L2TP LAC sessions of a maximum of 8K-1 different L2TP tunnel peers can be steered simultaneously within a single access router. [251198]

- Multi-Chassis Synchronization (MCS) and SRRP between a redundant BNG pair is only supported with a difference of up to two major SR OS software releases. The period in which different SR OS releases are deployed should be kept to a minimum, and Major ISSU is recommended to upgrade the earlier SR OS.
- DSCP remarking on access-egress is not supported for LAC, L2-aware NAT and sessions with a GTP uplink.
- When access-egress MTU is exceeded, no “IPv6 Packet Too Big” or “IPv4 Datagram Too Big” messages are sent for hosts with a GTP uplink. IPv4 packets without the DF bit set are fragmented.
- Externally-assigned subscriber identification strings should avoid using the ten character format “*xxxxxxxxxx*”, where *x* can be either “A-Z” or “0-9”. This format can overlap with the internally-generated subscriber identification that is enabled by default. For example, PPP or DHCP setup can fail, and debug will display a “Non auto-generated sub-id *xxxxxxxxxx* with an auto sub-id format not allowed” message. See TA 17-0884 for more information. [255628]
- VPRN of **type spoke** is not supported for subscriber management interfaces.
- Subscriber Access Bonding does not support Gx, Gy or NASREQ.
- PADI authentication is not supported for PPPoE sessions that should be bonded.
- Traffic steering of L2TP LAC is supported on a chassis populated with FP3-based line cards. Creation of a **steering-profile** is blocked if there are one or more line cards with FP2 or older, and the provisioning of line cards with FP2 or older is blocked if a **steering-profile** is configured.
- For traffic steering of L2TP LAC, downstream L2TP data traffic redirection to a VAS network next-hop IP address is only supported over an R-VPLS interface. If not on a R-VPLS interface, the L2TP session with a **steering-profile** goes to a steering-failure state with a log message.
- The following features are not supported on an IOM4-e-HS:
  - Access Node Control Protocol Management (ANCP)
  - Web Authentication Protocol (WPP)
  - PCC rule-based subscriber services (Gx or RADIUS)

11.55 PW-SAP for Epipe VLL Services

- Capture SAPs are not supported
- Ethernet CFM is not supported on PW ports or PW-SAPs.
- PW ports only support dot1q or QinQ encapsulation.

- The Independent Mode of PW Redundancy is not supported. That is, the PW port only acts as a slave from the perspective of PW preferential forwarding status.

11.56 VLL Spoke Switching

- If the control word is modified on a T-PE device in a pseudowire switched environment with either a Cisco or an Nokia router running a previous software revision as the S-PE device, it may be necessary to toggle the spoke binding status on the S-PE device (l2vfi connection in the case of a Cisco). [57494]

11.57 VPLS

- Remote MAC Aging does not work correctly due to ECMP, LAG or multiple paths that span different IOMs/IMMs/XCMs. If you have ECMP, LAG or multiple LSPs and a remote MAC learned on a given IOM/IMM/XCM moves to another IOM/IMM/XCM, the MAC will be first aged out of the FDB table when the remote age timer expires, even if the MAC is not idle. It will be then relearned on the new IOM/IMM/XCM. [33575]
- In a distributed VPLS configured with SDPs transported by MPLS (LDP/RSVP) where the ingress network interface for a given SDP is moving due to network events from one IOM/IMM/XCM to another IOM/IMM/XCM, the MAC addresses remotely learned on that SDP will start to age-out regardless of whether they are still active or not until twice their configured **remote-age** value is reached. Their ages will be then set back to 0 or the address will be removed from the FDB as appropriate. [47720]
- In a distributed VPLS configuration, it may take up to (2*(Max Age)-1) seconds to age a remote MAC address, and in cases of CPM or CFM switchover, it may take up to (3*(Max Age)-1) seconds. [48290]
- A user VPLS SAP might stop forwarding traffic after the SAP port bounces if that SAP is managed by a management VPLS (mVPLS) with Spanning Tree Protocol disabled. The workaround is to remove the mVPLS if the Spanning Tree Protocol is not required. If Spanning Tree Protocol is required, it should be enabled on the mVPLS. [60262]
- When a CPM or CFM switchover occurs during STP convergence, a temporary traffic loop or a few seconds of traffic loss may occur. [77948, 78202]

- The RSTP and MSTP Spanning Tree Protocols operate within the context of a VPLS or mVPLS service instance. The software allows for the configuration of an STP instance per VPLS service instance. The number of STP instances per VPLS or mVPLS service instance depends on 1) the number of SAPs/SDPs per VPLS and 2) the number of MAC addresses active within a VPLS.
- When using Ethernet Ring Automatic Protection Switching (R-APS) as defined in G.8032, CCMs and G.8032 R-APS messages continue to be forwarded in the control VPLS even if the service or its SAPs are administratively shut down. The Ethernet ring instance can be shut down to stop the operation of the ring on a given node.
- Provider-tunnels are not supported on BGP-AD R-VPLS Services.
- Per-service hashing will not work for egress VPLS management IP traffic in a VPLS service. [91377]
- LACP Tunneling for VPLS applies to untagged LACP only.
- ECMP and weighted ECMP are only supported for unicast traffic. The SR OS router will only select a single path for broadcast, unknown, and multicast traffic.

11.58 Routed VPLS

- Multicast traffic is incorrectly dropped if all of the following conditions are met.
  - A Routed-VPLS (R-VPLS) service is configured to allow the forwarding of IPv4/IPv6 multicast traffic from the VPLS to the IP side of the service.
  - Multicast traffic enters a SAP or mesh-/spoke-SDP in the VPLS side of the service which should be forwarded to a different SAP or mesh-/spoke-SDP in that VPLS service based on IGMP/MLD snooping state.
  - The shortest path to the source is across the R-VPLS interface. [209900]
- If PIM is configured on the IP interface of a routed I-VPLS service, any IPv4 multicast traffic sent over that interface will be flooded into the I-VPLS but not into the B-VPLS. [212347]

11.59 Proxy-ARP/ND

- Proxy-ARP/ND are not supported on the following services or in combination with the following features:
  - B-VPLS
  - I-VPLS

  - M-VPLS
  - R-VPLS
  - E-Tree
  - Subscriber-management, ARP-reply-agent, subscriber host connectivity verification (SHCV), residential split-horizon-groups, DHCP/DHCPv6, ARP-MSAP trigger, ARP-host configured.
  - VPLS Interface (although configurable, Proxy-ARP/ND is not supported) [220190]

11.60 IES

- In the saved configuration for IES services, the IES instance and interfaces will appear twice: once for creation purposes and once with all of the configuration details. This allows configuration items such as DHCP server configuration to reference another IES interface without errors. [56086]
- If two IES interfaces are connected back-to-back through a 2-way spoke-SDP connection with SDPs that have keepalive enabled and IGP is enabled on the IES interface with a lower metric as the network interfaces, the related SDPs will bounce due to SDP keepalive failure. The GRE-encapsulated SDP-ping reply will be ignored when it is received on an IES interface. [68963]

11.61 VPRN/2547

- VPRN service traffic with the DF (Do Not Fragment) flag set and requiring fragmentation to be transported through an SDP tunnel is correctly discarded, but an ICMP Type 3 Code 4 (fragmentation needed and DF set) message is not issued. [18869]
- The service operational state of a VPRN might be displayed incorrectly as Up during its configuration while some mandatory parameters to bring it up have yet to be set. [31055]
- Dynamic Multipath changes might not work in the case of VPN-IPv4 routes and might require a restart of the service. [31280]
- Each MP-BGP route has only one copy in the MP-BGP RIB, even if that route is used by multiple VRFs. Each MP-BGP route has system-wide BGP attributes and these attributes (preference) can not be set to different values in different VRFs by means of **vrf-import** policies. [34205]

- The **triggered-policy** feature does not apply to **vrf-import** and **vrf-export** policies in a VPRN. One needs to reset the target VRF instance in order to re-evaluate these policies or to disable the **triggered-policy** feature. [43006]
- Executing a **ping** from a VPRN without a configured loopback address may fail with a “no route to destination” error message despite there being a valid route in the routing table. The error message is misleading and should state that the reason for the failure is not having a source address configured. [55343]
- Misconfiguring the network so that two VPRNs leak the same prefix from VPRN to GRT results in only one leaked route in the GRT. After correcting the misconfiguration, an additional **shutdown** and **no shutdown** of the VPRN is required. [92147]
- VPRNs auto-bound to GRE tunnels cannot co-exist with IGP shortcuts since the line cards or CFM cannot forward GRE-encapsulated traffic for tunneled next-hops. [91863]
- Only regular IPv4 and IPv6 route-type routes leaked from the VPRN into the Global Routing Table (GRT) are supported. Unsupported route types are: aggregate, BGP-VPN extranet 6-over-4 IPv6, or 6PE IPv6 routes.
- If a VPRN is configured with **auto-bind-tunnel** using GRE and the BGP next-hop of a VPN route matches a static black-hole route, all traffic matching that VPN route will be black-holed even if the static black-hole route is later removed. Similarly, if a static black-hole route is added after **auto-bind-tunnel** GRE has been enabled, the blackholing of traffic will not be performed optimally. In general, static black-hole routes that match VPN route next-hops should be configured first, before the **auto-bind-tunnel** GRE command is applied. [167012]
- In case of multiple VPRNs on the same node when two VPRN routes with same RDs are compared, the VPN next-hop metric is used, which can be derived from either of the VPRNs. This causes inconsistent behavior when ECMP is enabled in one of the VPRNs. Toggling the operational state of one of the VPRNs can change the order of which route is selected. [197655]
- An SDP is always preferred over **auto-bind-tunnel** irrespective of the Tunnel-Table Manager (TTM) preference. [199763]

11.62 VRRP/SRRP

- The MAC address displayed for an SRRP gateway IP in the **show router arp** output on a subscriber interface does not show the MAC address of the Virtual Router but is that of the interface. Use the **show srrp** command to see the VR MAC address actually in use. [57838]

- If the **in-use** priority on each side of an SRRP connection goes to zero, both routers will incorrectly elect themselves as master. [60032]
- Under a VRRP policy, host-unreachable events can be configured. If the address configured is not reachable on the active CPM/CFM, the policy will use the configured priority to affect VRRP instances. Upon a High-Availability switchover, the address will be deemed reachable for a while. This period depends on the Interval and Drop Count configured under the event. Once the period is over, the policy event will properly reflect whether the address is reachable or not. [161154]
- After walking the tVrrpOpOperGroupName MIB object with SNMP using invalid indexes, successive walks on the object with valid indexes return no values. [249034]
- When the VRRP-aware PIM feature is configured in the Base router and a VPRN instance, state changes in the **oper-group** are not reflected in the VPRN. [252389]

11.63 VXLAN

- VXLAN R-VPLS services can only be bound to VPRN interfaces and not IES interfaces. [173106]
- When a BGP-EVPN route advertised from a Data Center (DC) controller has a VTEP endpoint (next-hop in the BGP-MH NLRI) in the same local subnet as the DC-PE’s egress network interface, the IP next-hop will not be resolved. It is required to have a Layer-3 router between the DC-PE’s egress network interface and the remote VTEP, or a /32 static route to the remote VTEP. [182672]
- The following limitations must be considered when using non-system IPv4 or IPv6 VXLAN termination in a VPLS/R-VPLS/Epipe service.
  - Assisted-Replication is supported on services using non-system IPv4 VXLAN termination but not on services using IPv6 VXLAN termination.
  - Ethernet Segment Identifier PBR/PBF is not supported.
  - IGMP-snooping is not supported.
  - When terminating VXLAN tunnels, the router does NOT check if there is a local Base router loopback interface with a subnet corresponding to the VXLAN tunnel termination address.
- Assisted-Replication has the following limitations.
  - Assisted-Replication leaf and replicator functions are mutually exclusive within the same VPLS service.

    - Assisted-Replication is supported along with IPv4 non-system-IP VXLAN termination; however, the configured **assisted-replication-ip** (AR-IP) must be different than the tunnel termination IP address.
    - The AR-IP address must be a /32 loopback interface on the Base router.
    - Assisted-Replication is only supported in EVPN-VXLAN services (VPLS with **bgp-evpn vxlan** enabled). Services with a combination of EVPN-MPLS and EVPN-VXLAN are supported; however, the assisted-replication configuration is only relevant to VXLAN.
- IPv4 VXLAN destinations on R-VPLS services or IPv6 VXLAN on Epipe/VPLS/R-VPLS do not support QinQ network-port encapsulation, since the maximum supported egress encapsulation is otherwise exceeded. When **config**>**system**>**ip**>**allow-qinq-network-interface** is executed, the configuration of R-VPLS services with VXLAN and IPv4 source VTEPs or Epipe/VPLS/R-VPLS services with IPv6 source VTEPs will not be allowed. Prior to Release 15.0.R6, the **allow-qinq-network-interface** CLI command was allowed but the BGP next-hop for EVPN-VXLAN routes in the above services were not resolved. When upgrading to Release 15.0.R6, **allow-qinq-network-interface** must be removed from the configuration if the VXLAN services mentioned earlier are configured. [257756, 266545]
- Network interconnect VXLAN Interconnect ESs (I-ESs) are supported on dual BGP-instance services. The following features are not supported on dual BGP-instance services with I-ES:
    - Proxy-ARP/ND
    - IGMP and PIM snooping
    - Assisted-Replication with leaf configuration
    - spoke-SDPs
    - BGP-MH sites

11.64 EVPN for VXLAN

- A given <VTEP, Egress VNI> pair is restricted to one given VPLS service; hence, a MAC route with the same <VTEP, Egress VNI> cannot be imported into two different services even if they have the same import-RT. The MAC will only be installed in one service. A trap will be raised to warn the user when there has been an attempt to add the same <VTEP, Egress VNI> to more than one service.

- The system IP-address is used in EVPN-VXLAN as the source VTEP of all the VXLAN packets and as the BGP next-hop in all the BGP-EVPN advertisements. When changing the system address, an administrative toggling (**shutdown**/**no shutdown**) is required in the BGP-EVPN context of the VPLS services so that the new system address is used as the BGP next-hop. Note that the system address cannot be changed as long as BGP-EVPN is administratively enabled (protected by CLI). The source VTEP of the VXLAN packets is changed immediately though, without any additional action [167775].
- In general, no SR OS-generated control packets will be sent to the VXLAN auto-bindings, except for ARP, VRRP, ping, BFD and CFM.
- Although xSTP can be configured in BGP-EVPN services, BPDUs will not be sent over the VXLAN bindings. BGP-EVPN is blocked in mVPLS services, however a different mVPLS service can manage a SAP/spoke-SDP in a BGP-EVPN-enabled service.
- **mac-protect** and **provider-tunnel** is not supported in EVPN-enabled VPLS services for VXLAN tunnels.
- **mac-move**, **disable-learning** and other FDB-related tools only work for data plane learned MAC addresses and therefore, not for control plane learned MAC addresses in EVPN-enabled services.
- VPRN interfaces bound to EVPN-enabled R-VPLS services do not support the following parameters: **arp-populate**, **authentication-policy**.
- BFD is not supported on EVPN-tunnel interfaces.
- EVPN-VXLAN BGP routes are not imported if the BGP next-hops are resolved over a non-network interface, for instance, an IES interface.

11.65 IPsec

- In a multi-active tunnel group setup, ICMP pings to the tunnel’s local address may fail. [140341]
- BFD over IPv6 over IPsec is not supported.
- IPsec DHCP relay uses only the **gi-address** configuration found under the IPsec gateway and does not take into account **gi-address** and **src-ip-addr** configuration below other interfaces. [224586]

11.66 PBB

- For access multihoming over MPLS for PBB Epipes, the following features are not supported: PW switching, BGP-MH, network-domains, **mac-ping**, **mac-populate**, **mac-purge**, **mac-trace**, or support for RFC 3107, GRE and L2TPv3 tunneling.
- ISID-level shaping on a B-SAP is not performed for traffic entering a Routed I-VPLS service which is forwarded over a B-SAP configured with **encap-defined-qos**. In this case, the traffic uses the normal SAP queues on the B-SAP rather than those associated with the **encap-defined-qos**. [217774]

11.67 Video

- A sequence of configuration changes, multicast traffic start and set top box activity may lead to a mix up between the (*,G) and (S,G) records on the MS-ISA. Nokia recommends configuring PIM SSM to avoid the issue.

  This may result in a slow FCC or unrepaired packet loss. The **show video channel** command has two entries in that case: one for (*,G) and one for (S,G). The FCC/RET counters should step up on the (S,G) entry, not the (*,G). If the (*,G) FCC/RET counters increments, the workaround is to use the **clear router pim database** command to get out of the state. [82353]
- In normal operating conditions, the RTP-sequence numbers for a channel are increasing monotonically. An equipment failure upstream of the video-interface (such as rewrapper-issue, intentional reset of sequence numbers) may lead to a situation where this assumption no longer holds. The MS-ISA may, depending on the channel characteristics, take up to 10 minutes to resume proper operation if such an event should occur. [110872]
- For FCC/RET:
  - up to four video groups are supported per chassis
  - if a chassis contains only IOM3-XP/-B/-C, IMM, and ISM, a maximum of six ISAs can be supported.
- For Ad Insertion (ADI), the frequency of IDR frames in the network and ad streams must be less than one IDR frame every 1.3 seconds.

11.68 Mirroring/Lawful Intercept

- Simultaneous Filter Logging and Service Mirroring on egress is not supported. When simultaneously performing filter logging and service mirroring at egress, the service mirroring operation takes precedence over the filter logging operation.
- If a dot1q SAP is being mirrored on an IES interface, DHCP responses from the server to the DHCP clients are not mirrored. A workaround is to mirror the port instead of the SAP. [40339]
- A redundant remote mirror service destination is not supported for IP Mirrors (for example, a set of remote IP mirror destinations). The remote destination of an IP Mirror is a VPRN instance, and an endpoint cannot be configured in a VPRN service.
- Multi-chassis APS (MC-APS) groups cannot be used as the SAP for a redundant remote mirror destination service. APS cannot be used to connect the remote mirror destination 7750 SR nodes to a destination switch.
- OAM **vccv-ping** is not supported on mirror service spoke-SDPs (or ICBs in the case of PW Redundancy being used for redundant mirror services). This is primarily because mirror traffic is uni-directional.
- LI/Mirroring at the LAC for subscribers using MLPPPoX access is not supported. Nokia instead recommends LI at the LNS.
- LI at the LNS for MLPPPoX (oE/oA/oEoA) subscribers is only supported with a **mirror-dest** type of **ip-only**. No other **mirror-dest** types are supported for MLPPP subscribers at the LNS.
- If q-tagged traffic is mirrored to a mirror-destination SAP and the SAP has an egress QoS policy containing IP-based reclassification, the IP-based reclassification is ignored. [132504]
- NAT-based lawful interception criteria (that is, **configure li li-source** *x* **nat** ... in CLI) can not be configured/triggered/used via RADIUS with the exception of L2-Aware NAT subscribers.
- Mirroring services and Lawful Intercept (LI) are not supported with a Segment Routing tunnel when the tunnel is used in a BGP shortcut and in resolving a BGP unicast label route.
- Mirroring of packets using ingress label (**debug**>**mirror-source**>**ingress-label**) is not supported with the following Segment Routing (SR) tunnel types: SR-ISIS, SR-OSPF, and SR-TE. [224677]

11.69 L2TPv3 SDP

- The implementation of L2TPv3 for SDP transport does not support:
  - Any L2TPv3 control plane functionality
  - Support sequence numbering
  - Fragmentation and reassembly
  - Session ID configuration or validation
  - Authentication – the only authentication of tunnel payload is performed through validation of Source Address, Destination Address, and the ingress cookie
  - Service multiplexing – each SDP will transport one spoke-SDP

  Unless explicitly mentioned above, most pseudowire/Epipe features are not supported on L2TPv3 SDPs or spoke-SDP bindings, including but not limited to:
  - Layer-3 functionality
  - Pseudowire shaping
  - Ingress/egress QoS functionality
  - Pseudowire switching
  - Active/standby pseudowire services and inter-chassis backup
  - PBB
  - Application Assurance
  - Hash-label
  - PW Status signaling

  Operators expecting to deploy this feature set should contact their Nokia engineering support teams.

11.70 NAT

- Executing a **traceroute** from an inside NAT interface may result in an unexpected source IP address in the response packet when the max session limit is exceeded. [91154]
- There are some limitations to the functionality of the Application Layer Gateways (ALGs) in combination with NAT64 due to the way the ALG translations are done.

When translating inside-information into outside information, IPv6 addresses are translated into IPv4 addresses without any issues, but when an IPv4 addresses is received in the payload of an incoming message, this address will not be translated because it is a random outside address and not a NAT address. In the NAT44 case, this is not an issue because the inside host can connect to this address, but in the NAT64 case, the inside host cannot connect to an IPv4 host.

This has an impact on the possible scenarios involving the ALGs:

  - SIP—The connection information in a SIP message describes the IP addresses and ports to be used to connect to the other party of the call. From the perspective of a client behind a NAT64 gateway, his own IP address will be translated correctly, but the IP address received from the other side may be an IPv4 address and will not be translated into an IPv6 address. Thus, the NAT64-client will not be able to initiate a connection to the other client. If only one client is behind a NAT64 gateway, SIP-calls are still possible. When client A (IPv4) can connect to client B (NAT64), client B can use this connection to connect back to client A. If both clients are behind the NAT64 gateway (the same or different), both clients will receive each other's IPv4 outside addresses and no client will be able to start the connection.
  - RTSP—Connection information in an RTSP message describes the IP address and ports to be used by the client to receive the actual video/audio/ etc. traffic. If the client is behind the NAT64 gateway, the server will receive correctly translated connection information and the client will be able to receive the data sent out by the server. If the server is behind the NAT64 gateway, the server will not receive translated connection information and the server will not be able to send out the data to the client.
  - FTP—Some servers may abort the connection when they receive the wrong type of address according to their current connection.

- The **config aaa isa-radius-plcy radius-acct-server source-address-range** command depends on the number of maximum ISAs configured in all NAT-groups, including the ISAs that were removed before the node rebooted. For every ISA, a unique source address is used.
- L2-Aware NAT is typically used with DHCP-proxy where the IP-address assignment to the ESM subscriber-host is handled via RADIUS. In this application, the same IP address can be assigned to multiple subscriber-hosts. This allows for IP address sharing between subscriber-hosts, which is the main purpose of L2-Aware NAT.

  In cases where L2-Aware NAT is used with DHCP relay (instead of proxy) where the IP address is assigned directly by the DHCP server, the IP lease can be extended only by DHCP rebind messages that are broadcasted. Any attempt to renew the IP lease by unicast DHCP renew message will fail.

This issue should not be a problem since the DHCP protocol will switch to multicast DHCP rebind after a few failed attempts to renew the IP lease via a unicast DHCP renew message.

- Policy-based Routing (PBR) is not supported in conjunction with L2-Aware NAT. In cases where PBR is enabled for L2-Aware NAT, traffic will undergo NAT but PBR will not be executed.
- Static 1:1 NAT is not supported for L2-Aware NAT, DS-Lite or NAT64.
- L2-Aware NAT is not supported on the Retail service in a Wholesale/Retail Routed-CO model. Large-scale NAT can be used instead.
- All ingress traffic subject to NAT has to ingress on an IOM3-XP/-B/-C or higher if deterministic NAT is configured on the service and if multiple ISA cards are present in the **nat-group**. If this condition is not met, tmnxNatMdaDetectsLoadSharingErr error events will be generated and traffic ingressing older IOMs, subject to NAT, will be dropped. [150597]
- SAA does not support ICMP Echo-Request for L2-Aware NAT hosts.
- The following options in the **ping** command are not supported for L2-Aware NAT hosts:
  - DNS resolution for L2-Aware NAT subscriber
  - rapid ping
  - **interface**, **next-hop** and **bypass-routing** options, all of which are used to determine the outgoing path for ICMP Echo Request message. This is not compatible with ICMP Echo Request in L2-Aware NAT where the outgoing path is dictated by the ESM subscriber, which is instantiated in SR OS.
- For L2-Aware bypass:
  - Is mutually exclusive with vRGW
  - cannot be combined with other ISA redundancy mechanisms (such as active-active and active-standby)
  - can be used only with L2-Aware NAT. No other NAT mode (such as, LSN44 or DSLite NAT64) can be enabled in the same NAT group when L2-Aware bypass is configured.
  - Sharing of IP addresses assigned to hosts is not allowed between the ESM/ L2-Aware subscribers within a given inside routing context
  - Multi-chassis redundancy is not supported in conjunction with this feature (MCS is not supported in conjunction with L2-Aware NAT).
- IPv6 Firewall will not work in combination with HTTP Redirect. [261408]
- Service-chaining VXLAN traffic sent out towards the Data Center (DC) will have the DF-bit set to one and the identification field set to zero for all packets. [271907]

11.71 Virtual Residential Gateway

- Enabling Virtual Residential Gateway (vRGW) should only be considered in environments that do not utilize SAA. These functions contend for the same resources, although they are not directly mutually exclusive. When both the connectivity-check and SAA functionalities are configured simultaneously, there are accuracy and resource contention issues.
- Static IPv6 hosts are not supported in vRGW. It is, however, possible to provide a static IPv4 host with a SLAAC prefix and use IPoE-linking to automatically create an associated IPv6 host.
- Wholesale/Retail is currently not supported in vRGW.
- Subnets provisioned for BRG pool management must lie in a pre-configured L2-Aware NAT inside prefix. The dynamic range of a BRG pool may not contain the configured L2-Aware NAT inside IP address.
- On regular group interfaces, only a single BRG is supported per SAP.
- There is a maximum of one SLAAC prefix per BRG.
- Idle-timeout is based on SLA-profile instance, not per host. For hosts under the same BRG sharing an SLA-profile, it is not possible to detect early disconnect of a single host.
- All SLAAC hosts under a BRG sharing the same prefix will use a common forwarding context downstream. For predictable behavior, the same SLA-profile should be used for each SLAAC host. This does not apply to hosts within a residential IPv6 firewall context.
- For vRGW, IPoE session pre-authentication using LUDB can only be used to pick up a RADIUS authentication policy.
- The residential firewall only supports IPv6 packets with up to 64 bytes of known extension headers. Packets not conforming to this limit will be dropped.
- Portal redirect is not supported for IPv6 hosts using the residential firewall. [261408]
- When using the PPPoE client with default ingress QoS on the Epipe service SAP or SDP, traffic will not be load-balanced over ISAs. Ingress QoS either needs to use policers or enable shared queueing.
- PPPoE client for vRGW is not supported on WLAN-GW group interfaces.

11.72 Application Assurance

- When deleting an application or an application group, statistics for the current accounting interval will be lost. The workaround is to first remove all references to the application and application group thereby allowing the accounting intervals to occur, and then delete the application or application group.
- For an active flow, when an application assignment is changed in an **app-filter**, or an **app-group** assignment is changed in an application, the flow count for the associated protocol is doubled.
- All subscribers being serviced by an ISA card must be removed from the ISA (configured as **isa-aa** or **isa2-aa**) prior to removing the card from an "application-assurance-group". [77394]
- Application Assurance does not support traffic divert to/from R-VPLS services; this includes traffic divert for SAP or spoke-SDP interfaces in both R-VPLS and linked IES/VPRN services.
- Only ESM subscribers (both static and dynamic via DHCP/RADIUS) are supported in a Wholesale/Retail VPRN configuration.
- In a Wholesale/Retail configuration, AA is supported on the ESM subscribers or on the aggregate traffic SAP facing the retailer's network, but not on both.
- When creating new AA group partitions, unique partition ID values should be used across all groups.
- When creating AA policers, unique policer names should be used across all groups.
- If hosts for a single ESM subscriber are present in multiple service instances, simultaneous traffic in the separate service instances with the identical IP 5-tuple may be mis-classified by AA. [91809]
- If Cflowd export from AA exceeds the rate that the CPM/CFM can process, Cflowd packets may be silently discarded. [91811]
- AA Redundancy Protocol (AARP) does not support multicast traffic.
- At a 1 Gb/s rate, a single TCP session or UDP flow must have an average packet size greater than 250 bytes. If the average packet size is less than 250 bytes, fairness between sessions/flows cannot be guaranteed. [98658]
- AARP is not supported on the 7750 SR-c4.
- During the small period of time it takes to create a new Seen-IP subscriber, packets to or from that subscriber may be recorded as policy-bypass errors. These policy-bypass error packets are correctly forwarded but are neither classified nor recorded against the subscriber. [139622]
- AARP is not supported between 7750 SR-c12 and non-7750 SR-c12 chassis types.

- PCRF has to reinstall, using a RAR, any AA-usage monitoring AVPs after an IPoE session migration process of AA ESM Gx controlled subscribers is completed.
- AA features that modify packets, such as HTTP redirect, HTTP enrichment, TCP MSS Adjust, or DSCP Remarking, will not process GTP untunneled packets. [228575]
- Application Assurance supports divert to/from a PBB interface with the following exceptions:
  - a SAP config of <port>*x.y*
  - satellite ports
  - PXC ports
- vPort Hashing over Multiple Forwarding Complexes does not interoperate with AA capabilities tied to a specific subscriber. When using vPort hashing, the **adapt-qos** link mode is recommended on the access interface.
- AA captive HTTP-redirect cannot be used in a WLAN-GW ESM subscriber deployment where L2-Aware NAT is configured. Alternatives are to use a WLAN-GW DSM subscriber deployment with L2-Aware NAT, or a WLAN-GW ESM subscriber deployment with large-scale NAT. [260531]
- Application Assurance is only supported in LDP-over-RSVP network deployments in IES and VPRN services. [277553]

11.73 Cflowd

- Cflowd is not supported on subscriber SLAs.
- Persistency of the Cflowd Global **if-index** is not supported. [148012]
- With the greater performance of Cflowd on the 7950 XRS and 7750/7450 CPMs, it is possible to generate more collector-bound packets than the CPM management Ethernet port can forward. In cases where Cflowd is expected to handle a very high number of flows, it is suggested that all collectors are made to be reachable in-band.
- Cflowd sampling traffic ingressing or egressing a non-Ethernet SAP has limited support. For non-Ethernet SAPs, the encapsulation will only be reported as zero. [162360]
- While Cflowd can be configured under SAPs on a 7450 ESS platform, Cflowd processing is not supported on these platforms, except on 7450 ESS-7 or 7450 ESS-12 platforms with mixed mode enabled. [162472]
- When Cflowd sampling is performed at the egress interface, the ingress interface index is not known. As a result, the ingress interface index field will always be set to zero (0) in exported flow data.

- The Cflowd sampling process does not sample the following types of control plane traffic bound for the system CPMs/CFMs:
  - IS-IS protocol traffic
  - BFD over LAG link members (uBFD)
  - Layer-2 control traffic in IP interface

11.74 sFlow

- Starting with Release 13.0.R6, scale limits for sFlow will be enforced to avoid IOM resource exhaustion. If sFlow is enabled on a port with more than 50 SAPs or on an IOM with more than 1600 SAPs, sFlow will be administratively disabled. The number of SAPs must be reduced to an allowed limit prior to re-enabling sFlow on the associated port or IOM. Nokia recommends reducing the number of SAPs below these limits before upgrading to Releases 13.0.R6 and higher. [216190]
- sFlow is not supported for PW-SAPs. [217715]

11.75 BFD

- When an SRRP instance uses its own BFD, L3 MC-ring cannot be enabled. BFD may be enabled in subscriber SRRP or MC-ring, but not both. [73063]
- When using multi-hop BFD for BGP peering or BFD over other links with the ability to reroute such, as spoke-SDPs, the interval and multiplier values should be set to allow sufficient time for the underlying network to re-converge before the associated BFD session expires. A general rule of thumb should be that the expiration time (interval * multiplier) is three times the convergence time for the IGP network between the two endpoints of the BFD session.
- Multi-hop BFD currently does not support tunneled routes (for example, **ldp-shortcut**, **rsvp-shortcut**, or static route with **tunnel-next-hop**). [135994]
- BFD VCCV on a BGP VPWS or BGP VPLS service may not interoperate with third-party implementations that require a response to a VCCV-ping echo request message in order to maintain the corresponding BFD session. [184152]
- Rx/Tx message counters for BFD sessions are not retained with a CPM/CFM High-Availability switchover; it is expected they restart from zero after the switchover. [250631]

11.76 OAM

- Timestamping the SAA versions of Loopback and Linktrace are only applied by the sender node. The total time of delay for Loopback and Linktrace tests includes the packet processing time of the receiver node, which may be very inaccurate depending on the CPU load of the receiver node at the processing time. Accurate results can be gathered through the use of Y.1731 **two-way-delay**, which includes native time stamping and the removal of remote processing times. [87326]
- If a **mac-ping** or **mac-trace** request is sent with an unknown source MAC address and there are multiple SAPs, the user will see duplicated results because the request is flooded to each SAP and each SAP sends a reply to the request message. This is the expected behavior. [16298]
- The **oam vprn-ping** and **oam vprn-traceroute** commands for VPRN in a hub-and-spoke topology using hairpin routing do not work. If a hub-and-spoke topology is used, the spoke site must be associated with the hub VRF or the default route created must point to the hub site not a black-hole. If not, some sites will not be reachable from the spoke site.
- The **oam vprn-ping** and **oam vprn-traceroute** commands do not work in a hub-and-spoke network topology with the 7750 SR or 7450 ESS in mixed mode, or 7950 XRS as the Customer Edge (CE) hub. As a workaround, the 7750 SR or 7450 ESS in mixed mode, or 7950 XRS will send a control plane response from the hub to the requester Provider Edge (PE) to confirm connectivity to the hub PE.
- OAM DNS lookups are not working correctly if the full DNS name is not provided. [54239, 54689]
- An OAM Service Ping request for a VPRN service is always sent over the data plane (over the spoke-SDP) and not through the control plane. A VPRN Ping should be used to send a ping request using the control plane for a VPRN instance. [58479]
- ATM OAM F4 cells on a VPC Apipe service are always sent with a PTI equal to four for SEG cells and a PTI equal to five for end-to-end cells. [75052]
- Even if **source-mac** is specified when using **oam cpe-ping**, the resulting ARP request packet sent to the CPE device will still use the chassis base MAC address. [85034]
- E-LMI is not supported on LAG interfaces.
- In scaled scenarios, SAA ETH-CFM tests configured to run in continuous mode may experience some probe packet loss. [90784]

- When SAA ETH-CFM continuous tests are configured and CPM- or CFM-redundant system is configured for **redundancy synchronize boot-environment**, the SAA ETH-CFM tests may experience some probe packet loss upon switchover during the Boot Environment Synchronization stage. [92500]
- **ldp-treetrace**, **ping** and **traceroute** may not work properly during an LDP-FRR event until IGP has converged, if originated on the node experiencing the failure and traveling over the link being protected. [115907, 121716]
- An **lsp-trace** of an LDP FEC can return a “DSMappingMismatched” error in the presence of ECMP paths. This is because the ingress LER selects the first ECMP next-hop provided by the responding LSR for populating the Downstream Mapping (DSMAP) TLV in the **lsp-trace** packet for the next TTL value. If the LSR hashing the packet for the next TTL value chooses a different downstream path to forward the packet, the error is returned by that downstream node.
- In order to properly trace the single path of a FEC, the user must add the **path-destination** option and enter a specific 127/8 address to be used in the IP destination address field of the echo request packet and in the DSMAP TLV such that the control plane and the data plane at the hashing LSR will use the same downstream interface. In addition, the user can discover all ECMP paths via the use of the **ldp-treetrace** command and trace all paths of the FEC. [150970]
- The following OAM tool commands are not supported with BGP-AD VPLS spoke-SDP and PMSI, and with BGP-VPLS spoke-SDP: **mac-ping**, **mac-trace**, **mac-populate** with **flood** option, **mac-purge** with **flood** option, and **cpe-ping**. [152529]
- The ETH-CFM primary-VLAN function will not extract ETH-CFM PDUs on QinQ Ethernet SAPs that specify an outer tag (x) and a value of zero for inner tag (<port-id |lag-id>:x.0) on the 7950 XRS platform. This is also the case for all other SR OS routers that enable the **new-qinq-untagged-sap** option. [153841]
- **sdp-ping** and **sdp-mtu** are not supported with an P2MP spoke-SDP used as an I-PMSI in VPLS context.
- **p2mp-lsp-ping** is not supported with an RSVP P2MP LSP or an MLDP FEC used as an I-PMSI in VPLS context [154657].
- **p2mp-lsp-trace** is not supported with an RSVP P2MP LSP used as an I-PMSI in VPLS context. [154659]
- Operators who opt to change the default values for **dot1q-etype** or **qinq-etype** will not be able to use primary-VLAN functionality. [154756]
- PBB-Epipes configured with spoke-SDPs must not have the **fault-propagation** option configured under any MEP attached to a spoke-SDP. This is an unsupported configuration for PBB-Epipes using spoke-SDPs. [163737]

- When OAM is to be originated/terminated in a SAP context on a given LAG with **per-fp-sap-instance** enabled, Nokia recommends using, at minimum, a one-second interval timer. When scaling SAPs on LAG, even larger timer values may be required, especially on older hardware. Failure to do so may result in OAM sessions going down during LAG-member port status changes. [175261]
- Sub-second CCM MEPs may not transition to a defect state for possibly six seconds upon IOM reset on the 7750 SR-a4/a8 and 7750 SR-1e/2e/3e platforms. [209430]
- The following OAM tools are not supported with Segment Routing (SR) IS-IS or OSPF tunnels:
  - PW-level OAM tools: **vccv-ping** and **vccv-trace** are not supported for PW-switching
  - Service-level OAM: **svc-ping**, **cpe-ping**, **vprn-ping**, **vprn-trace**, **mac-ping**, **mac-trace**, **mac-purge**, and **mac-populate**
- CPE-ping ARP packets will not egress a SAP defined as a **connection-profile** when the request is generated from the local node. Specific to an Epipe service, a remote issue of the **cpe-ping** command will traverse the network, across supported connection types, and be transmitted out of the **connection-profile** SAP without the application of a VLAN from the **connection-profile** range. [227023]
- **oam vprn-ping** does not work for either static or dynamic ARP. [233988]
- The **oam vxlan-ping reply-mode udp** option uses the UDP port allocated by IANA for VXLAN GPE (Generic Protocol Extension for VXLAN). Therefore, **reply-mode udp** is not recommended in networks where VXLAN GPE is deployed.
- The **udp** option for the **oam vxlan-ping** [**reply-mode** {**overlay** | **udp**}] command is not supported when the VTEP source is anything other than an IPv4 system address; the **reply-mode overlay** option must be explicitly used with a non-system IP source. For example, if the VTEP source uses the **vxlan-src-vtep** option, the **vxlan-ping** response will be discarded if **reply-mode overlay** is not specified.
- **cpe-ping** responses may be received and processed within Epipe and VPLS services using a PBB-EVPN transport, even when the service is operationally or administratively down. [249487]
- When originated on a BGP IPv4 label route with ECMP, **lsp-trace** next-hops can only exercise a maximum of 64 next-hops. The next-hops are selected by going over the resolved next-hops beginning with the first BGP next-hop until 64 resolved next-hops are selected. A responder node also can report a maximum of 64 next-hops in the echo reply message using the same above rule to select

them. A consequence is that a subsequent echo request message for the next value of TTL can be sent to the incorrect LSR downstream of the responder node and will return an error (rc=5 DSMappingMismatched) if the number of ECMP resolved next-hops for the BGP IPv4 label router at the responder node is higher than 64.

- The mtrace2 feature introduced in Release 15.0.R6 is not available in Release 15.1.R1.

11.77 E-Tree

- When configuring **root-leaf-tag** SAPs, the **root-tag** VID or the **leaf-tag** VID cannot be zero. Therefore the following SAPs are not supported as **root-leaf-tag** SAPs:
  - SAPs on null-encapsulated ports (root-leaf-tag SAPs must be on dot1q- or QinQ-encapsulated ports)
  - sap :0 root-leaf-tag leaf-tag X
  - sap :X root-leaf-tag leaf-tag 0
  - sap :* root-leaf-tag leaf-tag X
  - sap :X.Y root-leaf-tag leaf-tag 0
  - sap :0.* root-leaf-tag leaf-tag X

  Where X and Y are any VID value different from zero or *. The following SAPs are however supported as root-leaf-tag SAPs:
  - sap :X.* root-leaf-tag leaf-tag Y
  - sap :X root-leaf-tag leaf-tag Y
  - sap :X.Y root-leaf-tag leaf-tag Z

  Where X, Y and Z are any VID value different from zero or *.
- **root-leaf-tag** SAP/SDP-bindings are only supported in VPLS E-Tree and not in EVPN E-Tree.
- **pw-path-id** is not allowed for SDP-bindings configured in VPLS E-Tree services. This is valid for **root-ac**, **leaf-ac** and **root-leaf-tag** SDP-bindings. Static PWs are fully supported, however.
- No SONET/SDH with BCP encapsulation is supported in VPLS E-Tree services.
- The following features are not supported in VPLS E-Tree services:
  - BGP-AD, and BGP-VPLS
  - M-VPLS
  - R-VPLS
  - GSMP

  - VXLAN
  - legacy OAM commands (**cpe-ping**, **mac-ping**, **mac-trace**, **mac-populate** and **mac-purge**)
  - provider-tunnel
  - BGP instance 2
  - spoke-SDPs with L2TPv3 SDPs
- The following features are not supported in VPLS E-Tree and EVPN E-Tree SAPs:
  - capture SAPs
  - **eth-tunnel** SAPs
  - **eth-ring** – E-Tree SAPs can be used as **eth-ring** data SAPs but control G.8032 traffic is not supported in VPLS E-Tree or EVPN E-Tree services.
- The following features are not supported in VPLS E-Tree SDP bindings:
  - **vlan-vc-tag** under an **sdp-bind** when it is configured as **root-leaf-tag**.
- Proxy-ARP/ND is supported in EVPN E-Tree services but not in VPLS E-Tree services.
- In a scaled scenario, and typically when a new BGP peer is added, there is a risk of having temporary **leaf-ac** to **leaf-ac** BUM traffic between EVPN-E-Tree PEs. If the ingress PE receives the egress PE's Inclusive Multicast route prior to the leaf ESI-label, BUM frames from the **leaf-ac** will be forwarded to the egress PE without the leaf ESI-label, preventing the egress PE from filtering egress traffic to **leaf-ac**s. The filtering will work once the egress PE's leaf ESI-label is received and programmed.

11.78 DNSSEC

- Full DNSSEC validating resolver is not supported.
- DNSSEC AD-bit validation is not executed during the boot phase.
- DNSSEC AD-bit validation is not supported for the WLAN-Gateway GTP interworking function.

11.79 OpenFlow

- ofp_match oxm IPv6-label encoding is aligned to four bytes, not three bytes, although only 20 bits are relevant.

- of1DecodeOxmTlvInt [ERR]: icmpv4_type field cannot be masked; it is rejected even if the mask is all one.
- The OXM value should be the same after applying the mask. If not, it is rejected. [166673]
- A CPM/CFM switchover causes the TCP connection with the OpenFlow controller to bounce. Flow states are preserved. [167252]
- OpenFlow controller is not informed when, due to an operational event or configuration change impacting OF programmed rule, the programmed flow table action for Layer-3 PBR actions or Layer-2 PBF action is changed to or from drop or forward. The exception to this issue is steering to RSVP-TE or MPLS TP LSP.
- Hybrid OpenFlow Switch (H-OFS) is enabled by deploying an IPv4/IPv6 ACL that:
  - embeds an OpenFlow switch instance
  - or chains to a system filter that embeds an OpenFlow switch instance

  The OpenFlow-enabling IPv4/IPv6 ACL filter is supported in the following contexts:
  - **config**>**router**>**if**>**ingress**>**filter**
  - **config**>**service**>**ies**>**if**>**sap**>**ingress**>**filter**
  - **config**>**service**>**ies**>**if**>**spoke-sdp**>**ingress**>**filter**
  - **config**>**service**>**vprn**>**if**>**sap**>**ingress**>**filter**
  - **config**>**service**>**vprn**>**if**>**spoke-sdp**>**ingress**>**filter**
  - **config**>**service**>**vprn**>**network**>**ingress**>**filter**
  - **config**>**service**>**vpls**>**sap**>**ingress**>**filter**
  - **config**>**service**>**vpls**>**mesh-sdp**>**ingress**>**filter**
  - **config**>**service**>**vpls**>**spoke-sdp**>**ingress**>**filter**

  Deploying an OpenFlow-enabling ACL in other contexts is not blocked and should not be done in production networks. [199550]
- PORT_STATS and PORT_DESC (multipart types 4 and 13) are available for SR-TE LSP, but the counters (tx_packets and tx_bytes) are 0.

11.80 NETCONF/YANG

- The SR OS <candidate> datastore has limited support aimed at early evaluation and laboratory testing. It should not be used in production networks.

- The NETCONF port is not configurable. NETCONF sessions are supported on TCP port 830 (as required in RFC 6242). NETCONF sessions received on other TCP ports (including 22) are not supported.
- Some configuration leafs are immutable and can only be configured when an object (such as a list entry or member) is created. The SR OS NETCONF server will reply with an error if immutable parameters are attempted to be changed. When using the <candidate> datastore, the error will occur when a <commit> operation is requested (not during the <edit-config>). The object must be deleted/removed and then re-created with the new value for the immutable parameter if it needs to be changed. An example of an immutable parameter is the “customer-id” for a service.
- The CLI **candidate edit exclusive** lock does not prevent a NETCONF session from doing a successful <lock> RPC. A NETCONF <lock> can be taken even when the CLI **candidate edit exclusive** lock is being held by a CLI session. A CLI **candidate commit** will be blocked if a NETCONF <lock> is taken, even if the CLI **edit candidate exclusive** lock was taken before the NETCONF <lock>.
- The following items apply to the use of the Base-R13 SR OS YANG modules and data model (XML namespaces urn:alcatel-lucent.com:sros:ns:yang:conf-*-r13).
    - The alu-conf-log-r13.yang module does not correctly model the keys of the event-control list. The event-number is not included as a key due to limitations of the underlying infrastructure in handling parameters that are optional keys in CLI (the “no” form of event-control does not require the event-number). NETCONF edit-config requests and get-config responses can correctly use the <event-number> tag as a key (to write and read event-control configuration) but the YANG module does not model it.
    - The “choice” and “must” statements are not supported so the mutual exclusivity of some YANG objects cannot be indicated in the models.
    - When using the Alcatel-Lucent Base-R13 SR OS YANG modules in an <edit-config> on the <running> datastore, an explicitly-defined “delete” operation on a key leaf, regardless of the existence of the key leaf, acts as a “merge” operation. [212204]
    - Base-R13 YANG modules (with the running datastore) are non-transactional. The XML configuration data in an <edit-config> is processed serially, and each line takes operational effect as it is processed. The request requires the same ordering and has the same dependencies as CLI. The **rollback save** and **rollback revert** operations are available via NETCONF to give partial transactionality ("all or nothing" type behavior) when using the Base-R13 modules. For transactional NETCONF behavior, it is recommended to use the Nokia modules with the candidate datastore. Here are some examples of required NETCONF client behavior when using the Base-R13 modules:
        - the NETCONF client must shut down objects before it deletes them

    - the NETCONF client must remove children before removing their parent (for example, delete SAPs before deleting the service)
    - the NETCONF client must order the XML correctly. For example, the creation of a SAP-ingress QoS policy must come first in the XML before that SAP-ingress policy is referenced by a new SAP object.
- The NETCONF Base-R13 implementation is tightly linked to the CLI infrastructure. That linkage results in NETCONF behavior that follows many CLI behaviors and constraints. Some examples are listed below.
    - Many CLI commands require several parameters to be specified at the same time. For example, in the **configure service ies 1 subscriber-interface** *CB_1* **dhcp gi-address** *192.168.10.1* **src-ip-addr** command, the user must specify an *ip-address* after the **src-ip-addr** keyword is specified. The NETCONF equivalent also requires that the associated XML tags are present together in an <edit-config> request. So the <ip-address> tag must be in the same request if the <src-ip-addr> tag is present.
    - Some CLI commands have parameters that are keywords where the simple absence or presence of the keyword indicates whether the parameter is configured or not; there is not a **no** form for these keywords. These keywords are modeled in YANG as boolean types, but it is the absence of the associated XML tag in a request that indicates that the keyword is disabled instead of specifying "false" for the tag. The NETCONF infrastructure converts a false value for a boolean parameter into a CLI **no** form, which causes an error because there is not a **no** form for the keyword. An example of this case is the **src-ip-addr** keyword. An XML request with <src-ip-addr>false</src-ip-addr> is converted to **no src-ip-addr** in CLI; however, **no src-ip-addr** is not valid as part of the **gi-address** command. To clear the **src-ip-addr**, a NETCONF request must specify the <ip-address> tag without including the <src-ip-addr> tag in the request.
- Due to tight coupling between CLI and Base-R13 modules, non-standard XML output occurs in a <get-config> response in several scenarios when using the Base-R13 modules. Some examples of the scenarios where this occurs are as follows:
    - A <get-config> response may return containers that are empty (such as <dns></dns>). These empty containers occur in the same places as CLI **info** (or **admin save** configuration files) that has empty CLI branches (such as **dns** immediately followed on the next line of output by **exit**). RFC 6020 (YANG) does allow these empty containers (see section 7.5.8 of RFC 6020) but some tools may complain about them.
    - Containers and objects are repeated in a <get-config> response in some cases. SR OS NETCONF <edit-config> requests and <get-config> responses for the <running/> datastore contain ordered

content layer objects. Dependencies between objects sometimes require a part of a container or object to be configured first and then the rest of the container or object can be configured later (perhaps after other parts of the configuration model have been specified).

    - The <shutdown> leaf is repeated within a container or object in some cases. This is done, for example, in filters (for example, inside <management-access-filter><ip-filter>) so that the filter is first operationally disabled (<shutdown>true</shutdown>), then updated, and then finally operationally-enabled (<shutdown>false</shutdown>).
    - Leaf-list parent nodes are repeated for each leaf-list entry in some cases (such as the <member> leaf-list under <configure><system><security><user>).

- The following items apply to the use of the Nokia SR OS YANG modules and data model (XML namespace urn:nokia.com:sros:ns:yang:sr:conf-* and state-*)
  - The Nokia SR OS YANG modules have limited support aimed at early evaluation and laboratory testing and should not be used in production networks. These YANG modules are expected to be updated in subsequent releases without adhering to all the module update rules specified in RFC 6020 Section 10. Changes will not be backwards compatible. Some likely changes include:
    - The namespaces of the modules may change with the adoption of recommendations in *draft-chen-netmod-enterprise-yang-namespace*.
    - Some identifiers (for example, leaf names) may change (mostly to clarify, improve consistency, or fix errors).
    - Some objects may change from leafs to containers to lists without changing names or following the deprecation/obsoletion guidelines in order to improve the structure of the module.
    - The richness of the models will be improved by more fully modeling existing data model constraints (for example, indicate valid ranges of parameters, patterns, mutual exclusivity via new “choice” constructs, various “must” conditions, etc).
    - Some leafs may change types (for example, from “string” to “integer”).
    - Some strings may change to leafrefs. This is primarily for references to objects that are not yet part of the subset of configuration that is supported in the Nokia SR OS YANG modules.
    - Larger modules may be sub-divided into several smaller modules for better modularity and clarity.
    - Some default values may change, be removed, or be added.
  - The Nokia SR OS YANG modules cover a subset of SR OS configuration and state data. See the Nokia SR OS YANG modules to explore what parts of the configuration and state data are modeled.

  - SR OS does not order requests if a reference to an object is not currently implemented as a leafref. This results in failure at commit time in some scenarios; for example, creating an object is processed after processing a reference to that object. The workaround is to use a series of individual transactions (multiple sequences commits); for example, commit creating the object then commit referencing it. [226067]
  - SR OS does not support logging of operator actions (who issued which commands) via NETCONF with the Nokia SR OS YANG modules.
  - A NETCONF <commit> operation may fail in some cases due to ordering issues.
  - The “must” statement is not supported. Some constraints (for example, dependencies between leafs) cannot be indicated in the models.

11.81 ISSU

- ISSU is not supported in Release 15.1.

11.82 Telemetry/gRPC

- The SR OS gRPC server uses a single configurable (using the **sgt-qos** command) QoS value for all Telemetry data. The SR OS gRPC server ignores the "QOSMarking" leaf (if provided) with a "subscribe" RPC.
- The SR OS gRPC server does not support "ON_CHANGE" telemetry subscriptions.
- The Nokia SR OS YANG modules cover a subset of SR OS state data. See the Nokia SR OS YANG modules distributed with the SR OS release to explore which parts of the state data model are supported.
- The CLI incorrectly allows **auto-config-save** to be configured in the **configure system grpc** context.
- The following gRPC operations are not supported (as defined by gnmi.proto v 0.4.0):
  - rpc Capabilities
  - rpc Get
  - rpc Set
- The gRPC server supports JSON encoding. All other encodings defined in gnmi.proto (version 0.4.0) are not supported.

- Under overload condition (meaning, the output queue with Notification messages exceeds internally defined threshold) all current Subscriptions are canceled.
- The Alias concept (as defined by gnmi.proto version 0.4.0) is not supported.
- The ONCE, POLL, and ON-CHANGE subscription modes (as defined by gnmi.proto version 0.4.0) are not supported.

11.83 Soft Reset

- Although the data plane interruption during a Soft Reset is minimized, there is a brief (non-zero) traffic interruption. Transit protocol packets can be affected by this interruption.
- In scaled configurations, the following protocols may experience interruptions in peering sessions during a Soft Reset on the 400G line cards (for example, 4-port 100 GE) when using the default protocol timers:
  - Broadcast IS-IS (point-to-point IS-IS is not impacted)
  - RSVP
  - P2MP LSPs
  - LDP (T-LDP is not impacted).

  Increasing the protocol timers in the configuration will prevent interruptions in the protocol peering sessions. BFD (which is not impacted by the Soft Reset traffic interruption) could be used in conjunction with larger protocol timers in order to have fast failure detection.
- If the far-end node of an Ethernet OAM (802.3ah) session is not an SR OS router with the support for the vendor-specific Grace TLV, then the Ethernet OAM sessions are interrupted briefly during a Soft Reset and will take down the associated port and protocols running on that port. Ethernet OAM grace is disabled at the system level by default and must be enabled prior to an ISSU in order to take advantage of this functionality (**config system ethernet efm-oam**).
- LLDP information is lost when a card is Soft Reset, but relearned once the Soft Reset is completed.
- LACP sessions (Link Aggregation Control Protocol – IEEE 802.3ax standard, formerly 802.3ad) using the default “fast” timers may briefly go down during a Soft Reset (dependent on card types and configuration). The LACP sessions will recover within a few seconds. LACP sessions using “slow” timers will not go down during a Soft Reset.

- If the far-end node of an Ethernet CFM (802.1ag CC) or Y.1731 session is not an SR OS router with the support for the proprietary SR OS ETH-VSM Grace, then the Ethernet CFM or Y.1731 sessions are interrupted during a Soft Reset. ITU-T Ethernet Defect (ETH-ED) can be used in place of the pre-standard SR OS ETH-VSM Grace. Without Grace support, configured intervals of less than one second will result in the sessions going down. Intervals of one second may cause the sessions to go down in some cases (dependent on other configuration). Sessions with intervals of 10 seconds or higher will not go down even without the Grace support.
- Soft Reset outage times may be higher than expected if one or more line cards are Soft Reset while the standby CPM is rebooting. [73285]
- The architecture of some IMM cards prevents the **hard-reset-unsupported-mdas** functionality from being used for a manual **clear card** during a Minor ISSU. In most software upgrade cases, these cards can simply be Soft Reset (without the need for the **hard-reset-unsupported-mdas**), but if there is a mandatory firmware update on these cards, then they must be hard reset. The **hard-reset-unsupported-mdas** option is blocked for the following IMM types: imm1-40gb-tun, imm5-10gb-xfp, imm1-100gb-cfp, imm12-10gb-sf+, imm3-40gb-qsfp, imm-1pac-fp3, and imm-2pac-fp3. [158482]

11.84 FlowSpec

- For FlowSpec routes, there is no support for next-hop resolution, interaction of router policies and FlowSpec route NLRI fields, or configurable **prefix-limit**.
- Installed validated FlowSpec routes do not disappear when next-hop disappears.

11.85 Accounting

- The **extended-service-ingress-egress** record accounting is designed only for lower-scale deployments that require extra information and is not available in other types of records.
- When **extended-service-ingress-egress** record is selected for an accounting policy, the minimum **collection-interval** must be 15 minutes. The total number of SAPs that use the new accounting record type must not exceed 2048. [142879]

11.86 WLAN-GW

- The distributed RADIUS proxy is only guaranteed to handle Access-Request packets of up to 1024 bytes. [221041, 241114]

12 Resolved Issues

The following sections describe specific technical issues that have been resolved in SR OS releases. See also Known Limitations, as some known issues may have been moved to that section. Resolved issues from Releases 15.0.R1 to 15.0.R6 also apply to Release 15.1. Refer to the most recent *SR OS 15.0 Release Notes* for the summary of resolved issues in Releases 15.0.R1 through 15.0.R6.

Although the issues are stated in present tense, they have all been resolved.

Note:

- Bracketed [ ] references are internal tracking numbers.
- Issues that were resolved in earlier releases, but which were not documented until the current release, are marked **[NEW]** and are documented in the section for the applicable release.
- Issues marked as MI might have had a minor impact but did not disturb network traffic.
- Issues marked as MA might have had a major impact on the network and might have disturbed traffic.
- Issues marked as CR were critical and might have had a significant amount of impact on the network.

12.1 Release 15.1.R3

12.1.1 System

- Connector-based SFF DDM log events are not suppressed when the connector breakout port is configured as administratively down and the parent connector is configured as administratively up. [276959-MI]

12.1.2 Routing

- Calculation of FIB utilization as displayed in the **show router fib** *slot-number* **summary** CLI context is not as accurate as it should be. As a result, the displayed FIB utilization value may be lower than the actual FIB utilization. [280808-MI]

12.1.3 BGP

- BGP next-hop resolution may incorrectly use a route which is not the longest match when the longest match is rejected by a **next-hop-resolution** policy. [278639-MI]

12.1.4 IGMP

- In rare cases, when changing an IES/VPRN interface, that has IGMP enabled, from a SAP binding to an R-VPLS interface binding, all line cards in the node may reset. A workaround is to shut down the IGMP interface prior to making the configuration changes. [252471-MA]

12.1.5 Services General

- When **trigger-packet** data is configured on a capture SAP, any incoming packet may trigger a data-triggered host creation irrespective of the administrative state of the capture SAP. [271790-MA]
- In Release 15.1, the **service-name** CLI command is replaced with the **name** parameter on the **create** line (see Changed or Deprecated Commands in Release 15.1.R1-1 for more information). When upgrading a node from a prior release, this change may cause an IES or VPRN R-VPLS interface to stay down. In this scenario, the VPLS *svc-id* referenced in the new **name** command may not match the VPLS *svc-id* in the old **service-name** command, keeping the R-VPLS interface operationally down. The workaround is to remove the R-VPLS interface configuration from the IES or VPRN service and then re-add it with **vpls** *svc-id* with the correct *svc-id*. [281853-MA]

12.1.6 Subscriber Management

- When using Local Address Assignment (LAA) in combination with IPoE-linking for provisioning SLAAC prefixes for vRGW hosts, a different unique /64 SLAAC prefix is advertised to every host, even if they are in the same home. Instead, there should be one unique /64 prefix per home. [260919-MI]

- Removal of a retail subscriber interface containing an **address** entry with the **gw-ip-address** being the same as the **address** entry IP will cause the **address** entry to remain stuck in the wholesale subscriber interface, and will not allow the configuration of the same entry in the wholesale service or any retail subscriber interface using the wholesale service. To prevent this from happening, the **address** entry must be explicitly removed before removing the retail subscriber interface. If the issue occurs, the wholesale subscriber interface must be removed, or the node must be reset to resolve it. [275087-MI]
- An incoming DHCP discover with option 50 for an existing DHCPv4 lease incorrectly generates a RADIUS Access-Request packet based on DHCP offer instead of DHCP discover. This is only seen when neither IPoE session nor **re-authentication** is present. [275214-MI]
- Subscriber management memory usage may increase significantly when either **arp-host** or subscriber host connectivity verification (SHCV) is enabled and a continuous stream of ARP traffic arrives at a high rate. The recommended workaround is to rate-limit incoming ARP traffic by using distributed CPU protection (DCP) policies in combination with CPU protection MAC monitoring. [273133, 280949-MA]

12.1.7 VRRP

- The system incorrectly allows users to delete a LAG port associated with a VRRP policy. If the configuration is saved in that state, it will fail to execute after a node reboot. [277557-MI]

12.1.8 IPsec

- In an MC-IPsec scenario, the MC-IPsec standby node may incorrectly send an ICMP-redirect message to the source of a packet which is shunted to the MC-IPsec master. This behavior occurs when a routing loop in created through misconfiguration on the MC-IPsec master node, which forwards traffic to the MC-IPsec standby node. [271885-MI]

12.1.9 Mirroring/Lawful Intercept

- Creation and deletion of a configuration under the **config li mirror-dest-reservation** CLI context is only available in LI logs and not sent to LI SNMP. [280096-MI]

12.1.10 WLAN-GW

- IPv6 data-triggered authentication is dropped when the DHCP node is disabled under the WLAN-GW VLAN tag range configuration. [274518-MI]

12.2 Release 15.1.R2

12.2.1 Hardware

- Prior to Release 15.1.R2, the MIB object tmnxHwTemperature used the value of -1°C to indicate that the hardware component does not contain a temperature sensor. Beginning in Release 15.1.R2, the value of -128°C indicates this. [277545-MI]

12.2.2 System

- When IEEE 1588 Port-Based Timestamping (PBT) is enabled on a port, and the port is **shutdown** and then re-enabled, IEEE 1588 message correction fields may become corrupted. To work around this issue, remove and reconfigure **ptp-hw-assist** on the router interface associated with the port, or remove and reconfigure the PTP port. [275145-MI]
- On the 7750 SR-1, if there is an active connector port alarm on an MDA, and the MDA is **shutdown** or removed, then the connector port alarm is not cleared or removed from the **show system alarms** output. [275941-MI]
- In Release 15.1.R2, all connector-based ports now have a minimum ingress-shaping rate of 10 Mbps, which can be increased in steps of 10 Mpbs. The ingress-shaping rate is configurable in the **configure port** *port* **ethernet ingress-rate** CLI context. [278377-MI]

12.2.3 Routing Policies

- Removing the last entry in the last policy-statement from **router policy-options** without performing **commit** or **abort** will prevent the standby CPM/CFM from synchronizing with the active CPM/CFM after a reset of the standby CPM/CFM. In this case, the standby CPM/CFM will stay in "syncing/standby" state indefinitely until a **commit** or **abort** command is executed. [275666-MI]

12.2.4 OSPF

- When an OSPF designated router (DR) learns various options advertised by adjacent routers, the DR correctly performs a logical OR function on those options and returns the result in its LSA responses. However, when the LSAs containing various options disappear from the OSPF database, the DR still keeps sending out the original options without recalculating the value resulting from the OR operation using the new options. [255986-MI]
- When translating from Type-7 to Type-5 LSAs with zero metric, the OSPF External Metric Type incorrectly changes from E2 to E1. A workaround is to not use a zero metric value. [274578-MI]

12.2.5 MPLS/RSVP

- When both interface statistics (**show router interface statistics** or **show router interface detail**) and MPLS statistics (**show router mpls interface statistics**) are accessed simultaneously, either via two CLI instances or via a CLI and an SNMP instance, a CPM/CFM High-Availability switchover may occur. [275959-MI]

12.2.6 Filter Policies

- A **filter log** configuration with a scaled number of **destination memory** entries may fail to execute after a node reboot due to the scaling limit being reached. [257560-MI]

12.2.7 BGP-EVPN

- Remote EVPN MAC/IP routes received with the same route-distinguisher (RD) as the BGP-EVPN service's local RD may lead to MACs becoming unresponsive in the FDB, even when the remote MAC/IP route has been withdrawn. VPLS service RDs should always be unique in the network. See RFC 7432 for more information. [252740-MI]

12.2.8 Services General

- When the XMPP server is configured to use a VPRN instead of the **management** routing context, the router may incorrectly use the DNS servers configured in the **bof** instead of the ones in the VPRN to resolve the XMPP server name. [278476-MI]

12.2.9 Subscriber Management

- In rare cases, when a RADIUS CoA message triggers an SLA-profile change on a redundant BNG system, a RADIUS Accounting-interim-update message (reason sla-stop) may be sent after an SRRP switchover instead of a RADIUS Accounting-interim-update message (reason SRRP switchover). [273308-MI]

12.2.10 IPsec

- In very rare cases, an IPsec gateway configured for EAP authentication with active IPsec clients may result in an unexpected CPM/CFM High-Availability switchover. [275422-MI]

12.2.11 Video

- Collecting an **admin tech-support** file may cause an ISA configured as **isa2-video** to reset. [276791-MA]

12.2.12 NAT

- Configuration of Static Port Forwards (SPFs) using a NAT-policy with an IP-filter associated will be saved in the wrong order. The SPF statement in the configuration file precedes the declaration of the NAT-policy in use by the IP-filter(s), which results in a failure to execute the configuration after a node reboot. Before loading such a configuration, the saved file must be manually modified to remove the SPFs and either put them after the declaration of the NAT-policy IP-filter(s), or execute the **config** command after a full configuration load. See TA 17-0988a for more information. [270593-MA]

12.2.13 WLAN-GW

- The following packets generated by a WLAN-GW do not indicate their lack of fragmentation support by setting the DF bit in their IP-header.
  - all types of soft-tunnels: GRE and L2TPv3
  - DHCPv4 offer, ACK, and NAK for migrant and DSM states
  - distributed RADIUS proxy and DSM RADIUS proxy: AAA Access-Request towards a AAA server
  - RADIUS client for migrant state authentication: AAA Access-Request towards a AAA server
  - distributed RADIUS proxy and DSM RADIUS proxy: AAA Access-Challenge towards AAA client
  - AAA Accounting-Request
  - AAA Accounting ON/OFF
  - CoA ACK/NAK for migrant or DSM states

  Also, the HTTP redirect performed in portal state towards the client is missing a valid IP identification field. [275093-MI]

12.3 Release 15.1.R1-1

12.3.1 Hardware

- In a 7750 SR-7 or 7450 ESS-7 chassis equipped with SF/CPM5, the **show system switch-fabric exclude-sfm 2** CLI command incorrectly indicates a minimum forwarding capacity of 50% for CPM-A and CPM-B. [271980-MI]

12.3.2 System

- During a CPM/CFM boot sequence, accounting files and log files stored on compact flash are deleted if they were created 12 or more hours ago. [269385-MI]
- Deconfiguring a syslog server address and assigning the address value again while the syslog is assigned to a *log-id*, can result in syslog packets being generated with a random UDP destination port value that is different from the explicitly configured value or the default port value of 514. A workaround is to remove and then re-add the *log-id* that uses the syslog. This ensures the correct UDP destination port is used again. [270474-MI]

12.3.3 NETCONF

- After a <copy-config> is performed with <startup/> as <target> and <url> as <source>, a **show bof** command displays that the primary-config points to the URL's configuration file instead of just replacing the contents of the primary-config's configuration file with the contents of the URL's configuration file. [231237-MI]

12.3.4 DHCP

- The **delegated-prefix-length** and its **minimum** or **maximum** within the **dhcp6 local-dhcp-server pool** CLI context can be incorrectly set to an invalid value of zero via SNMP. This results in a configuration that fails to execute after a node reboot. [269160-MI]

12.3.5 BGP

- When **update-fault-tolerance** is enabled followed by a CPM/CFM High-Availability switchover, BGP may interpret some attributes from the imported VPN routes as malformed. A second CPM/CFM High-Availability switchover may trigger unexpected BGP update messages leading to some traffic loss. [268016-MI]

12.3.6 LDP

- Adding and removing IPv6 addresses on an operationally-down interface using the **no ipv6** CLI command may result in an LDP unusual error message. [266146-MI]

12.3.7 QoS

- The default value displayed in the **queue drop-tail low percentage-reduction-from-mbs** CLI context has been changed from **10** to **default** in the following default QoS polices:
  - network-queue “default”
  - shared-queue “default”
  - shared-queue “policer-output-queues”
  - shared-queue "egress-pbr-ingress-queues"
  - queue group templates ingress
    - queue-group “_tmnx_nat_ing_q_grp”
    - queue-group "_tmnx_nat_ing_q_grp_v2"
    - queue-group "_tmnx_lns_esm_ing_q_grp"
  - queue group templates egress
    - queue-group "_tmnx_nat_egr_q_grp"
    - queue-group "_tmnx_nat_egr_q_grp_v2"

  This is a display change only. The default value used remains at **10**, so the operation itself is unchanged. [261593-MI]
- For multiple QoS PIR- and CIR-related MIB objects that have a Hi value for the upper 32 bits and a Lo value for the lower 32 bits, setting the Hi and the Lo values at the same time via SNMP incorrectly fails. [267685-MI]

12.3.8 Services General

- MACsec Key Agreement (MKA) protocol tracks peers via a Member Identifier (MI). There is a log event when the number of potential peers exceeds the configured value (MACSEC #2005). This may indicate an incorrect MI. [274602-MI]

12.3.9 Subscriber Management

- The maximum number of identical IPv4 or IPv6 framed routes with different next-hops that are installed in the routing table is determined by the **ecmp** *max-ecmp-routes* configuration in the routing instance. If there are more identical IPv4 or IPv6 framed routes in a routing instance, they are kept in the shadowed state based on the lowest next-hop IP address as the tie breaker. Keeping more than 512 identical IPv4 or IPv6 framed routes in shadowed state can lead to increased redundant CPM/CFM reconcile times and should be avoided. [222227-MI]

12.3.10 IPsec

- In rare cases, receiving multiple IKE-SA-INIT requests with different NAT ports during a remote-access tunnel setup, immediately followed by the system trying to delete those tunnels when DPD expires, may cause an unresponsive tunnel to remain in the IPsec gateway. An attempt to clear such tunnel may cause the ISA (configured as **isa-tunnel** or **isa2-tunnel**) to reset. [265376-MI]
- A remote-access tunnel may not be completely deleted from the system if:
  - the peer is behind a NAT router and NAT-T is configured
  - the IKE SA (phase-1) was rekeyed with a different UDP port than used to set up the tunnel initially as a result of the NAT router expiring the translation
  - the tunnel goes operationally-down because DPD expired on the SR OS node [266085-MA]

12.3.11 Cflowd

- The interface index value is incorrectly embedded in the IPv6 link-local next-hop address for IPv6 Cflowd flows. This interface index is displayed in the output of the **tools dump cflowd** CLI command and is also present in the next-hop address portion of the flow that is exported to the collector. [265707-MI]

12.3.12 BFD

- Multi-hop BFD sessions can bounce if the BFD packets are forwarded via a static route resolved by an RSVP-TE LSP (with the **static-route**'s **tunnel-next-hop** set to **rsvp-te**) that is protected by FRR when the LSP moves to the FRR protection path. [260103-MI]

13 Known Issues

Following are specific technical issues that exist in Release 15.1.R3 of SR OS. See also Known Limitations, as some known issues may have been moved to that section.

Note:

- Bracketed [ ] references are internal tracking numbers.
- Known Issues added in this release are marked **[NEW]**.
- Issues marked as MI have a minor impact and will not disturb network traffic.
- Issues marked as MA may have a major impact on the network and may disturb traffic.
- Issues marked as CR are critical and will have a significant amount of impact on the network.

13.1 Hardware

- The optics modules details displayed in the output of the **show port detail** CLI command may be displayed in hexadecimal notation instead of the normal decimal notation if the optics modules parameters were incorrectly programmed to include non-printable ASCII characters. The specific value is appended with “(hex)” to indicate such an occurrence. [84012-MI]
- Back-to-back runts may not be counted correctly under port statistics on 100GE ports. Also, some runts may be counted as fragments. [129447-MI]
- The system marks any IOMs/IMMs/XCMs as “failed” if they have rebooted due to an internal failure more than five times in a period shorter than or equal to 25 minutes. Marking the cards as “failed” and generating log messages is currently also done for the standby CPM. This is incorrect since the standby CPM cannot be prevented from rebooting. [149975-MI]
- FCS errors on received frames on a 100G Ethernet port may incorrectly cause the “ingress FCS errors” alarm to be reported against the ingress forwarding complex of the line card. This alarm should only be reported for FCS errors due to an internal defect. This issue is resolved for the x4-100g-cfp2, x4-100-cxp, imm4-100gb-cfp4, and imm4-100gb-cxp cards but may still occur on other cards with 100G Ethernet ports. [228977-MI]
- Ingress FCS errors on Ethernet Ports are incorrectly counted as Threshold Drops on the port. This issue is resolved for the x4-100g-cfp2, x4-100-cxp, imm4-100gb-cfp4, and imm4-100gb-cxp cards but may still occur on other cards with Ethernet ports. [229141-MA]

- After a Soft Reset on the p1-100g-tun-b card, the Maximum Rx Per-Channel Power field in Coherent Optical Port Statistics incorrectly displays 0.0 if the Maximum was a negative value before the reset. [231536-MI]
- Removing the active CCM while pressing the LT button will cause all LEDs to remain flashing in test mode on the other CCM and all XCMs/XMAs. Pressing the LT button for one second and releasing it clears the test mode. [232315-MI]
- When using copper SFPs with ma44-1gb-csfp or ma2-10gb-sfp+12-1gb-sfp MDA, late collisions are detected with an Ethernet configuration of half duplex and a speed of 100 Mbps. [253719-MA]
- With copper SFP (3HE11904AA) or cSFP (3HE10113AA), a configuration of half duplex and a speed of 100 Mbps is not supported. [253773-MA]
- On the p160-1gb-csfp IMM, ma44-1gb-csfp, or me40-1gb-csfp MDAs, if a cSFP in the bottom row (or left hand side if mounted vertically) is removed and reinserted and one of the two ports for that cSFP is configured as a sync-if-timing reference, then the sync-if-timing reference may go into an LOS state. If this occurs, disable and re-enable the sync-if-timing reference to recover it. [255081-MA]
- An IOM with an m2-oc192-xp-xfp MDA incorrectly counts ingress FCS errors when receiving packets of 54 bytes or less at line rate. If **fail-on-error** is configured, a burst of small packets may generate enough ingress FCS errors on a complex for the IOM to be disabled into failed state. [261195-MI]

13.2 Satellites

- If an Ethernet satellite is configured with an incorrect satellite type (that is, not matching the actual satellite), the satellite may fail to become active once this is corrected and may need to be manually rebooted. [223753-MI]
- Executing a satellite configuration file immediately following a CPM activity switch on the host may result in an SNMP set failure on the TDM satellite. This causes the satellite to reset. [251276-MA]
- Due to inconsistent CLI/SNMP checks, an operator may be able to delete an uplink binding while a client port served by that uplink still has a non-default **encap-type**. If an **admin save** is done at this time, the resulting configuration file will not successfully reload. The workaround is to change the **encap-type** back to **null** for all client ports before removing the associated port-topology uplink binding. [253541-MI]

13.3 CLI

- Special characters (“\s”, “\d”, “\w”) do not work with pipe/match functions. [100089-MI]
- Removing or adding certain candidate configuration can trigger false CLI warnings like "Deleting non-existing node ..." or "Referencing non-existing object ... ", while the candidate configuration change is valid and applied correctly. [226091-MI]
- A CLI **rollback revert** operation fails toward a configuration with **maximum-routes** *value* or **maximum-ipv6-routes** *value* defined. [252516-MI]

13.4 System

- The system incorrectly allows an **admin save** operation initiated by a user to be aborted if another user initiates another **admin save** from another session. [79185-MI]
- If no new events are logged after the retention period, a file will not be created on the compact flash card. A CLI **show** of the **log-id** will then give a false error: “MINOR: CLI Could not access”. [94600-MI]
- Copying a file to a TFTP destination sometimes prompts for a confirmation to overwrite the destination file on the TFTP server, even if that file does not exist. [120649-MI]
- A CLI **rollback revert** operation that requires the change of certain attributes on channels that are associated with a channelized SONET/SDH ports may shut down the base port in instances where the shutdown is not required. [121080-MI]
- CPU-protection policies are not supported at the IES/VPRN tunnel-interface SAP-level/context but in some cases, it is incorrectly shown as configurable. Note that a CPU-protection policy (if desired) should be applied at the tunnel-interface level instead of at the tunnel-interface SAP-level. [133148-MI]
- Traffic load balancing is less efficient when the number of BGP next-hops to a prefix is greater than eight and where the number of resolving links for some BGP next-hops is greater than eight as well. [198707-MA]
- An SSH or Telnet session that has **login-control ttl-security** enabled, and hence has a per-peer-queue created, currently does not display the per-peer-queue in the output of the CLI command **show system security per-peer-queuing detail**. [241794-MI]

- Rebooting a node after upgrading the switch fabric cards from SFM4 to SFM5 without first provisioning the new **card-type** correctly will result in a **card-type** mismatch system alarm. The alarm will remain even after the cards are provisioned correctly. A High-Availability switchover will clear the alarm. [244620-MI]
- If the port of a mirror destination is manually **shutdown**, and there is a large number of mirror sources associated with the destination, it may take some time for the **shutdown** to complete. During this time, the administrative and operational states may not match. [249531-MI]
- A benign message may appear on the console of a 7950 XRS CPM card when booting: "B:sysMonitor*pri0:COMMON:tmPcieSwitchGetBdbErrorEquivalent suppressing pcie error for bitmap: 0000000010000000" [251539-MI]
- The “Time of last boot” in the **show card detail** output might be incorrect. If a card (for example, CPM) comes up with an old system time from its onboard clock, and that system time is then later corrected by some means (for example, by synchronizing with an NTP server), the “Time of last boot” might be incorrect. [252267-MI]
- When a port on an me2-100gb-cfp4 or me2-100gb-qsfp28 MDA is used as SyncE reference into the central clock and an LOS condition on this port is detected, the central clock will switch to another reference if available. During this switch, a phase transient that exceeds the limit defined by the standards may be observed. [253138-MI]
- On the ma44-1gb-csfp or me40-1gb-csfp MDAs, if a cSFP in the bottom row (or left hand side if mounted vertically) is removed and reinserted, then IEEE 1588 Port-Based Timestamping may no longer function properly for the two ports. If this occurs, disable and then re-enable the PTP Ethernet ports or remove and re-add **ptp-hw-assist** to the router interfaces using those ports. [255211-MI]
- In some cases, when the system is handling a large number of SNMP-GET requests, a Major ISSU may take longer time than expected to complete. [256056-MI]
- In the **config li mirror-dest-template** CLI context, **router** "**management**" is incorrectly listed as an option in the CLI, but is not supported. [258975-MI]
- The **config li mirror-dest-template** CLI context only supports **layer-3-encap ip-udp-shim** encapsulation. The encapsulation type **ip-gre** is incorrectly listed as an option in the CLI. [263419-MI]
- Synchronous Ethernet (SyncE) will not be enabled if an MDA with SyncE configured is rebooted while it is administratively down. The Transmit timing selected and Sync interface timing status will be missing from **show mda detail**. This issue is present in Releases 14.0.R4 and higher. A workaround is to reconfigure SyncE while the MDA is administratively up. [274333-MI] **[NEW]**

- On the 7750 SR-1, if PTP over IPv4 with loopback interfaces (such as **system**) is used, then a source-address must be configured within the **configure system security source-address application ptp** *interface-name* CLI context to avoid corruption of the correction-field of PTP packets. [275252-MI]
- Unusual error events may occur when RADIUS sends invalid match string information in the supplied authorization profile (for example, unknown command). [275562-MI]

13.5 ATM

- When a non-terminating ATM SAP (**atm-vpc** or N:1 connection-profile) is implemented on a multi-chassis-APS (MC-APS) group, and both MC-APS member ports fail, the SAP will source ATM ETE-AIS cells onto the pseudowire, in addition to setting the lacIngressFault and lacEgressFault pseudowire status bits. The opposite SAP, at the other end of the pseudowire, will send out the AIS cells, while also generating its own in response to the PW status change. This results in the opposite SAP sending AIS cells at a rate of two per second instead of one. There are no false alarms or other ill effects, and both AIS cell flows stop when service is restored. [147334-MI]
- The option to set the CLP bit to 1 in the ATM cell header for traffic egressing the non-expedited queues for an IES or VPRN service is not supported on IOM3 or higher. If the functionality is enabled via the **configure qos atm-td-profile** *td-profile-id* **clp-tagging** on the ATM traffic descriptor assigned on SAP-egress for a IES/VPRN service, and that SAP resides on an MDA which is on an IOM3 or higher, there will be no tagging of the ATM cells corresponding to the traffic from non-expedited queues. [235800-MI]
- In an ATM **connection-profile**, it is possible to configure the members using PVP encapsulation values. An ATM **connection-profile** should only support PVC values in vpi/vci format. This can also be done in SNMP, but this action is blocked in the CLI. [243704-MI]

13.6 SNMP Infrastructure

- The system may not correctly count the number of failed SNMPv3 authentication attempts in the event-control log. [64537-MI]
- SNMP replay events may not function properly for replay functionality with multiple trap-targets pointing to the same address (even if they belong to different trap-groups/logs). This issue does not affect replay functionality with only one trap-target per trap-receiver address. [69819-MI]

- The system may not return a lexicographically higher OID than the requested OID in an SNMP GET-NEXT operation when incorrect values are used. This behavior is seen in the tcpConnectionTable table. [80594-MI]
- After 497 days, any “Last Change” counter on the system will wrap around due to a 32-bit timestamp limitation. The “Last Oper Chg” value in the output of the **show router interface** command is one example of such counter, but there are numerous other cases where this limitation applies. [83801-MI]
- A system that does not have a system IP address or a management IP address configured may not be able to generate SNMP traps. [98479-MI]
- SNMP traps are not forwarded when overwriting or modifying existing trap-target in both the base and VPRN context. [177129-MI]

13.7 LAG

- The **weight-threshold** option is not supported in combination with the **standby-signaling power-off** command. This invalid configuration combination is permitted by the CLI, and does not work as expected to bring down the LAG when configured. These two options should not be configured together. [241334-MI]
- In certain scenarios, where some ports in a LAG with **per-fp-sap-instance** enabled have either been removed or are down, service self-generated-traffic (sgt) packets sent over a spoke-interface using that LAG may be incorrectly dropped. [280912-MI] **[NEW]**

13.8 MLPPP

- If an MLPPP bundle with more than one link has **magic-number** configured and all links are looped back, a link may not become active when it stops being in a looped-back state. To recover from this and to allow the link to become active, shut down the bundle and toggle the **magic-number** attribute. [143509-MI]

13.9 APS

- Individual APS channel group members may be reported as down while the APS port status is operationally up. This is strictly a display issue. [89341-MI]

- If a CLI **rollback** operation must remove or alter the working bundle associated with a BPGrp, then it will also delete and rebuild any APS port associated with that BPGrp. [121024-MI]
- A CLI **rollback** operation that requires the removal of member links from a multilink bundle or BPGrp will shut down the associated bundle or BPGrp during the course of its operations, even if one or more member links still remain throughout the course of the rollback. [121066-MI]
- If all APS ports are active on either the working or protect router with a highly-scaled MC-APS configuration including MLPPP BPGrps and that router reboots, some PPP links may suffer PPP keepalive failures during the APS switchover process. In that case, the link will bounce and renegotiation will occur. [156523-MI]

13.10 ATM IMA

- When an IMA group is deleted while the group still contains IMA member links, some of the member links may show erroneous DS1 and DS0 ingress statistics after the deletion. [151573-MI]

13.11 Routing

- When no management IP address is defined on a node, **ping** *IPv4-address* **bypass-routing** does not work. [245940-MI]
- In a routing policy edit, the use of square brackets, including an enumeration of more than two digits, immediately followed by curly braces in the same regular expression, may result in a CPM/CFM reset. [281169-MI] **[NEW]**

13.12 Routing Policies

- In a route-policy entry, a match on a community logical expression, composed of community sets and logical operators, will not take into account the exact keyword and associated logic attached to any community set that is itself a logical expression. [274021-MI]

13.13 IPv6

- Enabling the **advertise-tunnel-link** command in IGP while both IPv4 and IPv6 IGP routes are present can result in IPv6 routes on neighboring nodes to end up with the wrong metric, which may cause non-optimal routing or routing loops. IPv6 routes do not support the **advertise-tunnel-link** option and their metric should not be affected. The workaround is not to enable the command when IPv6 routes are present. [247162-MA]

13.14 DHCP

- An IP address that is released and immediately granted again by the master **local-dhcp-server** may, in rare cases, result in a false positive alarm on the standby failover **local-dhcp-server**: “BNDUPD message could not be processed for DHCP lease * – reason: hostConflict”. [177704-MI]

13.15 IP/RTM

- IPv6 packets resolved over IPv4 IGP-shortcuts have the TTL of their MPLS label set to 255 instead of inheriting the IPv6-header's hop-limit value. Packets with an explicit null label are not affected. This behavior applies to both CPM-/CFM-originated and transit packets and cannot be changed when the user toggles the CLI for the TTL propagation over IGP-shortcut for either type of packets. [254050-MI]
- When FRR is activated at the ingress LER of an SR-TE tunnel, CPM-/CFM-orginated packets over the SR-TE LSP are dropped. [264579-MA]

13.16 IS-IS

- When used in combination with ECMP, the **show router isis lfa-coverage** command may provide incorrect results. [142527-MI]
- The IS-IS **lsp-refresh-interval** cannot be set to the previous value if configured after the **lsp-lifetime**; it will always be set to 50% of the **lsp-lifetime** value. The workaround is to set the **lsp-refresh-interval** to a different value before setting the **lsp-lifetime**, and then setting the **lsp-refresh-interval** to the desired value. [231950-MI]

- IS-IS adjacencies over hybrid QinQ network ports may bounce after a CPM/ CFM High-Availability switchover. A workaround is to use at least the default **hello-interval** of 9 seconds. [249071-MI]
- In certain scenarios, after an RSVP interface using IS-IS traffic-engineering has bounced due to a port flap, it may incorrectly not be considered for CSPF calculation. A workaround to recover is to bounce the port or interface again. [280088-MI] **[NEW]**

13.17 BGP

- Changing the BGP **router-id** value in a base or VPRN configuration will immediately cause a flap of all BGP neighbors that are part of that instance. [121246-MI]
- The CLI **show** command for MVPN BGP routes does not correctly filter on **originator-ip**, **source-ip**, and **group-ip** addresses. This is the case when filtering with the default addresses in MVPN-IPv4 and with any MVPN-IPv6 addresses when no **type** is given. [185058-MI]
- The following FlowSpec NLRI subcomponent type 12 "fragment" values are mapped to wrong filter match criteria:
  - [e=1 a=0 len=1 not=0 m=0/1; LF=0 FF=0 IsF=0 DF=0] => fragment false
  - [e=1 a=0 len=1 not=1 m=0/1; LF=0 FF=0 IsF=0 DF=0] => fragment true

  The correct behavior would be:
  - [e=1 a=0 len=1 not=0 m=0/1; LF=0 FF=0 IsF=0 DF=0] => no fragment/ fragment off (any packet matches)
  - [e=1 a=0 len=1 not=1 m=0/1; LF=0 FF=0 IsF=0 DF=0] => no packet matches [208414-MA]
- While upgrading to Release 14.0.R4 or higher, when **advertise-external** and **add-path** are enabled on an ASBR that is also a route reflector, the device may not be configured properly. When a route from an I-BGP peer is best, no best external path is injected, and the internal route is not reflected to other clients. Nokia recommends removing **advertise-external** from the configuration of the device and relying on **add-path** for path diversity. See TA 17-0919a for more information. [228990-MI]
- The number of extended communities displayed in the **show router bgp routes flow-ipv4** and **show router bgp routes flow-ipv6** CLI commands is limited to ten. [262669-MI]
- The **show router bgp routes mcast-vpn-ipv6 hunt all** command does not display Internal or Local routes. [263221-MI]

- A value 86400 cannot be configured for **advertised-stale-time** under the **long-lived family** *family* CLI context. This is considered as configuring a default value for the parameter and the value specified under **long-lived** is still inherited by the family. [263802-MI]
- If the ASBR does not have any VPRNs and ORF capability negotiated, an End-of-RIB message for the **vpn-ipv4** family is not sent by the RR to the ASBR. [263803-MI]
- In some cases, when a PE has been configured to send BGP ORF updates for extended-communities, upon changing policies affecting route-target extended communities which are used to import BGP-VPN routes into local VPRNs, ORF updates may be sent by the PE with the incorrect route-target list. This may affect how routes are advertised by the receiving PE or a route-reflector, and how the received routes are imported into local VPRNs. [276450-MA]
- BGP next-hops for 6PE and labeled-IPv6 routes are incorrectly resolved by less-specific black-hole static routes configured on the router, despite more specific routes being present in IGP. This causes the 6PE and labeled-IPv6 routes learned from the BGP next-hop to be black-holed as well. The issue is present in Releases 15.0.R4 and higher. A workaround is to not use any black-hole static route which can resolve the BGP next-hop. [282249-MA] **[NEW]**

13.18 BGP-EVPN

- When an EVPN route is withdrawn because of a parsing error, a withdrawn log event is generated but an additional log event to indicate the reason of the error is not always generated. In some error cases related to an EVPN mpunreach attribute, there is no log event generated when this attribute is ignored. [184549-MI]
- The use of the same import route-target for multiple VPLS services is not currently recommended. [205726-MI]
- In a scenario with EVPN all-active multihoming, a PE part of the Ethernet Segment (ES) may learn a MAC “M1” in the FDB associated to a local SAP in the ES, but as type **evpn** (this is possible if “M1” was learned on a peer ES PE and subsequently advertised in EVPN). In this situation, new frames received on the local ES SAP with MAC SA = M1 will not trigger the relearning of M1 as type “Learned”, as would be expected. This may generate some unnecessary extra flooding from remote PEs if the peer ES PE withdraws M1. [208989-MA]
- More than one EVPN P2MP leaf to the same root node within the same VPLS service may result in system instability. [270965-MA]

- When a VXLAN all-active Interconnect ES (I-ES) is used, performing a **shutdown**/**no shutdown** on the I-ES does not flush the existing MAC addresses that belong to the VXLAN network and are associated to MPLS destinations in the FDB. In this situation, the router may black-hole traffic from the MPLS network to any of those MACs, even if there is a valid MAC/IP route coming from the VXLAN network. A workaround is to **shutdown**/**no shutdown** a **bgp-evpn vxlan** or **bgp-evpn mpls** to resolve the issue. [280160-MA] **[NEW]**

13.19 MPLS/RSVP

- A non-CSPF LSP path whose next-hop is over an unnumbered interface will not come up if traffic engineering is disabled in IS-IS or . In addition, RSVP needs the router ID of the next-hop to look up an existing neighbor or to create a new neighbor before sending out the PATH message to the local and remote borrowed interface address. This information is looked up in the Traffic Engineering (TE) database. [146593-MI]
- For LSPs over unnumbered interfaces, routed messages such as RESV, RESVTEAR and PATHERROR are destined to the remote-router ID. A successful RTM lookup for the packet destination is necessary to send the message. If the IGP is shut down, then RTM lookup will fail, and the message may get dropped. [153707-MI]
- When using an unnumbered IP interface as a Traffic Engineering (TE) link for the signaling of RSVP P2P LSP and P2MP LSP, it is required that all nodes in the network have their **router-id** set to the system interface. [153791-MI]
- Under certain conditions and topology, there is a chance that a one-to-one detour originating from a PLR will be incorrectly merged by a detour merge point such that the detour terminates back onto the same PLR. [157528-MI]
- With unnumbered RSVP interfaces, the RESV message from an LSR to its upstream neighbor can use a different interface than the PATH message. If the authentication parameters of the links used by the PATH and RESV messages are different, either they use a different key, or authentication is disabled in one of the links; the upstream LSR detects the authentication mismatch and discards the RESV message. The LSP will not come up.

  The reason is that the RESV packet is actually routed to the upstream neighbor. This is not an issue with numbered interface since the upstream neighbor uses the local interface address in the Previous Hop (PHOP) object in the PATH message and thus, the RESV is always routed via the link used by the PATH message and representing the same subnet. With unnumbered interface, the PHOP object uses a loopback address of the upstream neighbor that corresponds to the borrowed IP address of the unnumbered interface used by the PATH message. Thus, routing back to this loopback address can use a different link than the one used by the PATH message which does not necessarily follow the shortest path due to CSPF. It can also be due to asymmetric routing over the link, and this issue will occur even if the PATH message used the shortest path.

The workaround is to configure the same authentication parameters on all RSVP interfaces, numbered or unnumbered, where a RSVP packet may be sent or received. [160106-MI]

- All TIMETRA-MPLS-MIB TimeInterval objects over 248.5 days and using a TimeInterval of TIMETRA-TC-MIB (for example, vRtrMplsLspTimeUp) returned negative values. This issue has been resolved for most objects, excepting those still using the TimeInterval format: for example, vRtrMplsLspPathStatTable and vRtrMplsP2mpInstStatTable. [223032, 229059-MI]

13.20 LDP

- LDP Path-MTU Discovery is not reducing the Path MTU correctly in presence of IGP-shortcuts if the MTU of the tunnel is less than the MTU of the interface at the ingress LER. [140723-MI]
- Modifying the system-interface IP address may cause LDP to keep the old IP address in the LIB/LFIB as a local prefix binding. To remove this binding, the LDP’s administrative state must be toggled. [149930-MI]
- When transitioning from a peerTemplate-driven T-LDP session to a manually-configured T-LDP session with **local-lsr-id** enabled, the session will flap. [165590-MI]
- As part of the Auto T-LDP feature, peerTemplates are saved in the configuration file based on the order of creation. When a **rollback save** is performed and subsequently the user deletes or recreates the same peerTemplate, thus altering the template creation time, the **rollback revert** operation is not capable of reverting the template configuration based on the initial creation order at the time of the **rollback save**. [166160-MI]
- When an LDP IPv6 sub-interface is configured in a native IPv4 system (that is, one that does not contain any IPv6 configuration), the session flap causes the LDP sync timer to start once the adjacency comes up. It is not terminated even after receiving all of the End-of LIB LDP messages (prefix IPv4, prefix IPv6, P2MP IPv4 and P2MP IPv6) for the IPv4 session. The timer continues to its configured expiry time while the IPv6 session is operationally down. [224489-MI]

- On the 7750 SR, the supported TCP encryption algorithms for LDP and BGP sessions are **aes-128-cmac-96** and **hmac-sha-1-96**. As these encryption methods have been implemented as pre-standard Internet-Draft (I-D) and are not fully compliant with RFC 5926, they are not interoperable with third-party vendor routers. [236922-MI]

13.21 IGMP

- Padded IGMPv2 query messages received on a Routed-VPLS interface are incorrectly interpreted as IGMPv3. [281625-MI] **[NEW]**

13.22 PIM

- In rare cases, interfaces may have the same IPv6 link-local address, which is used as the primary interface address for IPv6 PIM. If the interfaces in the RP tree and shortest-path tree have the same IPv6 link-local address, then the router will be unable to send RTP-prune messages. [152125-MI]
- **lag-usage-optimization** is supported only when per-flow, MID-based hashing is enabled on a LAG and when no queue or SAP optimizations are enabled on the LAG. The configuration is not blocked when the condition is not met, and using **lag-usage-optimization** may lead to disruptions in multicast traffic. [180482-MI]
- In some cases, the "Curr Fwding Rate" in the output of **show router pim group detail** may incorrectly show a value after traffic for this multicast group has stopped. [202141-MI]
- Shutting down and deleting an interface rapidly (for example, using a script) may cause some multicast traffic not to be forwarded to other interfaces that are part of the Outgoing Interface lists (OIF lists) containing the deleted interface. To prevent this from happening, the interface should be deleted at least five seconds after it becomes operationally down. To recover from the incorrect state, the affected multicast groups can be toggled with the **clear router pim database** command. [203559-MA]
- When the Route-Distinguisher (RD) is changed without shutting down the VPRN, MVPN routes may not generate the source-AD routes with a new RD. [232158-MI]

13.23 QoS

- Egress policed packets redirected to an egress port queue group in a criteria action statement will always use the default queue group instance configured within the SAP egress QoS policy under the SAP. Consequently, any VXLAN VNI queue group redirection to a different queue group instance as part of an applied **queue-group-redirect-list** will be ignored. [243559-MA]
- When multiple service classes have an aggregate rate applied on a High-Scale QoS IOM (IOM4-e-HS), a small amount of priority leakage can occur where lower-priority classes forward packets rather than higher-priority classes. This can be alleviated by shaping each higher-priority scheduling class to a rate below the aggregate rate or setting the **low-burst-max-class** to be the highest low-priority scheduling class. [254865, 257370-MA]
- Setting a user-defined CPM queue's **mbs** and **cbs** value to a maximum of 131MB is incorrectly allowed. [271667-MI]
- Executing the **clear service id** *svc-id* **sap** *sap-id* **queue-depth** command may, in very rare cases, cause the active CPM/CFM to become unresponsive. [249752, 281598-MA] **[NEW]**

13.24 Filter Policies

- When removing a filter that has a **default-action deny** from a SAP or interface, a very small number of packets may be dropped. [92351-MI]
- If the ingress or egress ACL/QoS filter entry resources on any line card are close to full utilization (above 90% of capacity) for a given filter type, then the performance of some configuration updates to these filters may be degraded, especially during large configuration changes when using long filter match-lists, or large embedded filters. Configuration update performance degradation does not impact data-path performance of the line card. [161389-MI]
- On the 7750 SR-1e/2e/3e and 7750 SR-a4/a8 chassis types, when the CPM-filter entry number is equal to or greater than 1538, the forward or drop statistics per entry incorrectly always remains zero (0). CPM-filter logging and CPM-filter statistics on filter entry numbers in the range from 1 up to 1537 function correctly. To avoid this issue, CPM-filter entries can be renumbered. [269386-MI]
- On 7750 SR-1,SR-1e/2e/3e, SR-a4/a8, and SR-c4/c12 platforms, a line card reset may occur when CPM queue statistics retrieval and CPM queue creation happen at the same time. [280605-MA] **[NEW]**
- Applying FlowSpec filters to LI may result in missing filter entries if the entry ordering is changed from the server. [281054-MI] **[NEW]**

13.25 Services General

- A combination of **control-word** and **force-qinq-vc-forwarding** should not be used in a VPLS service when **proxy-arp** is enabled. This will lead to the ARP flooded frames being malformed. [222071-MA]
- The **vc-id** used for mesh-SDPs is by default the **service-id**. If the **vc-id** (which can be the **service-id** or a different value) is specified when creating the first mesh-SDP, then all subsequent mesh-SDPs use the same value. Specifying a different value will cause an error. [265483-MI]
- AGI TLV does not display the VPLS-ID in the output of **show router ldp bindings service-id** *service-id* **detail** when using BGP-AD for LDP-signaled pseudowires. [266992-MI]
- In some cases, when SPB is used with B-VPLS, "protected-mac" traps/event logs may be generated every 600 seconds. This occurs when IS-IS multicast frames are received on VPLS ports which are not on the path to the multicast source. These events are informational and do not impact service. [268575-MI]
- MACsec Key Agreement (MKA) protocol negotiates the active key server among the peers in the session. Multiple key server switches, when the key server priority is changed manually, may result in a previously-used key identifier being distributed by a non-Nokia key server. This results in a MACsec encrypted traffic outage. A workaround is to perform a MACsec **shutdown**/**no shutdown** on the affected port. [274712-MA]
- When enabling MACsec for a port, if the **ca-name** is configured while the port is **shutdown**, and before enabling **connectivity-association**, then MKA for the port will be in a failed state. Once in this state, a workaround is to remove and re-configure the **ca-name**. To avoid the issue, configure the **ca-name** after the port is **no shutdown**. [275131-MA]
- An R-VPLS interface is not supported in a VPRN **type spoke** service. [277035-MA]

13.26 Subscriber Management

- A DHCP ACK returned by a VPLS DHCP proxy will be incorrectly tagged and not reach the DHCP client in case the VPLS SAP where the client connects to is not a service delimiting tag or the outer customer tag. [147457-MA]
- SCTP source or destination port ranges match in IPv4 and IPv6 ingress or egress subscriber management credit control filters is not supported. [199371-MI]

- Gx Usage Monitoring is not supported in a dual-homed configuration. The error reporting via Error Message AVP to indicate this not-supported behavior is missing for Gx PCC-Rules. [211556-MI]
- When a subscriber *a* comes up with **accu-stats** enabled and the subscriber is renamed to *b* using the **tools perform subscriber-mgmt re-ident-sub** *a* **to** *b* CLI command, and there is a previously existing inactive subscriber *b* with offline statistics, then the previously stored offline statistics of subscriber *b* are overwritten. This behavior is seen only for renaming using the **re-ident-sub** command. [252148-MI]
- Data-triggered host setup is not supported in a Wholesale/Retail configuration when the retailer has **private-retail-subnets** configured. [269987-MA]
- In a stateful multi-chassis redundant setup, subscriber hosts created through data-trigger can fail to be synchronized on standby when data-trigger is disabled again. The error logged in such a case is: "VSA Alc-Force-DHCP-Relay received for a host not located on a ESM group interface, on which data-trigger hosts are enabled". [280659-MI] **[NEW]**
- When a **category-map** is unconfigured, the system incorrectly deletes any categories that exist in this **category-map** automatically. However, any **exhausted-credit-service-level** ingress or egress IP-filter or IPv6-filter entries that exist for these categories, that should also be automatically removed, are not removed. Because the filter entries without a category are not actually-installed ACL filter entries, there is no immediate traffic impact. However, these hidden entries will re-appear if the same **category-map** and categories are re-configured. A workaround is to delete all categories before deleting a **category-map**. [281601-MI] **[NEW]**
- Toggling the **private-retail-subnets** option on a retail subscriber interface after it has been configured as unnumbered prevents connectivity to subscriber hosts instantiated on that retail subscriber interface. [282073-MI] **[NEW]**

13.27 VPLS

- In a VPLS using an I-PMSI and a spoke-SDP of **vc-type vlan**, when L2PT or BPDU-translation is enabled on the service and STP BPDUs are received over P2MP leaf, they are incorrectly dropped as “Bad BPDUs”. [134168-MI]

13.28 VRRP

- IPv6 using VRRPv2 is not supported. IPv6 requires VRRPv3. If an IPv6 VRRPv2 advertisement is received, a log event is incorrectly raised. Statistics for invalid version messages should instead be counted and displayed using the **show router vrrp statistics** CLI command. [263708-MI]

13.29 VXLAN

- When a VXLAN tunnel is terminated in a VPRN service (instead of the Base router), a VXLAN egress VTEP **oper-group** may not go down when the VTEP route is no longer active in the VPRN routing table. [266540-MI]

13.30 IPsec

- When using IKEv1 tunnels, the system expects to receive the DPD capability in the first message from the peer. Some vendors send the DPD in the second message, which results in DPD functionality being disabled on the SR OS node, but enabled on the peer side. In the absence of ESP traffic, DPD may time out on the peer side and take down the tunnel. [280966-MI] **[NEW]**

13.31 Video

- In some cases, clearing the video interface statistics can cause it to incorrectly show a higher “Tx FCC Replies” count than the “Rx FCC Requests” count. [182951-MI]
- In rare cases when using a multicast-service, adding a new primary MS-ISA to an existing video group may cause some FCC/RET requests and multicast traffic to not be forwarded to all MS-ISAs in the group. The recovery action is to re-provision the affected MS-ISAs. [189479-MA]

13.32 WLAN-GW

- If subscriber-management persistency is enabled, WiFi UE mobility between access points (APs) can fail in some cases, displaying the following drop reason in DHCP debug traces: “Problem: There is currently another transaction active for this lease state”. The workaround is to disable subscriber-management persistency. [195056-MI]
- The Call Trace **live-output** option for WLAN-GW DSM UEs will only forward packets in the management router and any router or VRF where a **nat outside** context or DSM IPv6 pool manager is configured. [224993-MI]

13.33 NAT

- Dynamic ports are always reserved, even if only deterministic port blocks have been reserved via configuration. [195357-MI]
- With **nat-group** *nat-group-id* **redundancy active-active** configured, during full reboot or after **no shutdown** of the **nat-group**, traffic may be loaded on the first ISA’s coming up and revert to a stable, balanced load over all ISAs in the group shortly thereafter. [210575-MA]
- On scaled configurations with many static port forward entries present, some ISA cards may take a very long time to become active after a node reboot. [215131-MA]
- Traffic counters in the **nat inside downstream-ip-filter** are zero after a CPM/CFM switchover. [235940-MA]
- 1:1 L2-Aware subscriber hosts are not supported in combination with routed-CPE (**anti-spoof nh-mac** under the SAP). L2-Aware 1:1 can only work with one host IP, but in case of **nh-mac**, there could be more routes/addresses behind the subscriber. [240130-MA]
- Changing both **classic-lsn-max-subscriber-limit** and **dslite-max-subscriber-limit** while there are NAT policies associated with a router instance can, in rare cases, result in service impact and an unexpected error event: "BB_MGMT:UNUSUAL_ERROR Slot *x*: natPortRangeAddDetMap: maskedBits *y* is inconsistent with V4HashPrefLen ...". An ISA card reset is required to restore service. [270929-MA]
- DNAT-only flows are not logged with IPFIX. [271540-MI]
- Certain downstream IPv4 ICMP packets may result in an ISA reset when they are translated into an IPv6-packet by the NAT64 function on the ISA. [279862-MA] **[NEW]**

13.34 MSDP

- Logs may incorrectly show an MSDP peer transitioning from established to a lower state when the remote peer has not been configured to accept MSDP sessions and has a higher IP address. This does not cause any service impact. [161762-MI]

13.35 Application Assurance

- Under unexpected fragmented GREv1 traffic conditions, benign trace errors may be seen. [212589-MI]
- Under unexpected asymmetrical traffic conditions, sessions may be in analysis longer than expected, resulting in the delayed application of non-default policy and reporting of statistics. [271208-MA]

13.36 BFD

- Upon reset of an ASAP MDA, IS-IS may not re-register as a BFD client on multilink bundles. [62885-MI]

13.37 OAM

- **oam vprn-trace** packets incorrectly time out when sent to ASBRs in an inter-AS configuration. [59395-MI]
- Executing **oam host-connectivity-verify subscriber** *sub-ident-string* **sla-profile** *sla-profile-name*, will not trigger an ARP Request or Neighbor Solicitation if, as *sla-profile-name*, the *default-sla-profile-name* is used. [140038-MI]
- A reply to a **p2mp-lsp-ping** of an MLDP FEC will fail at the leaf LSR if the latter is enabled with the multicast upstream FRR feature (**mcast-upstream-frr** option) and has activated LFA next-hop towards the backup upstream LSR. [162937-MI]

- When a port member in a LAG changes from a non-operational to operational state, a sub-second CCM-enabled QinQ Tunnel Facility LAG MEP (LAG + VLAN) associated with that LAG will experience a timeout condition which will cause attached services to propagate fault. It is suggested that these Facility MEPs use a minimum CCM interval timer of one second. [175240, 200980-MI]
- If ETH-CFM is configured on a SAP in a BGP-EVPN VPLS where **cfm-mac-advertisement** is enabled, and the MAC used for the MEP/MIP is the SAP’s physical port MAC, when the card goes offline, EVPN will not withdraw the MAC route corresponding to MEP/MIP MAC. As a workaround, a specific MAC address for ETH-CFM in the BGP-EVPN VPLS service SAPs can be configured. [213818-MI]
- For option B inter-AS **p2mp-lsp-ping ldp** *p2mp-identifier* **vpn-recursive-fec**, the root system IP-address is not present in the leaf RTM. As such, the echo reply is forwarded via the unicast VPRN infrastructure. The system IP-address of the leaf node has to be configured within the VPRN and advertised via MP-BGP to the root node. Only IPv4 is supported as the advertised system IP-address; currently there is no support for IPv6 system IP-addresses. As such, for VPNs that only support IPv6 (that is, 6VPE), for **p2mp-lsp-ping ldp** *p2mp-identifier* **vpn-recursive-fec** to work, the user must configure an IPv4 system IP-address within the VPN. [230342-MA]
- **p2mp-lsp-ping ldp** *p2mp-identifier* is not supported in inter-AS option C scenarios in which the EVPN PE uses basic opaque FEC to resolve a root IP address existing in a remote AS. [243327-MI]
- MPLS **shortcut-local-ttl-propogate** and **shortcut-transit-ttl-propagate** for SR-TE LSP shortcuts are unsupported. [262159-MA]
- An OAM timestamp may drift away from the system clock after several days of uptime. A workaround is to perform a CPM High-Availability switchover to temporarily resolve the issue. This issue only affects 7750 SR-a4/a8 and 7750 SR-1e/2e/3e chassis. [274739-MI]

14 Change History for Release 15.1 Release Notes

The following table lists significant documentation changes to the SR OS 15.1 Software Release Notes.

Table 32 **Change History**

Part number	Date of Issue	Reason for Issue and Changes to Documentation
3HE 13648 0003 TQZZA 01	March 2018	Third 15.1 Release Notes.
3HE 13648 0002 TQZZA 01	January 2018	Second 15.1 Release Notes.
3HE 13648 0001 TQZZA 01	December 2017	First 15.1 Release Notes.

Customer Document and Product Support

Customer Documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation Feedback

Customer Documentation Feedback

3HE 13648 0003 TQZZA 01