Service Router | Release 19.10.R6

SR OS Software Release Notes

3HE 15407 0010 TQZZA 01

Issue: 01

June 2020

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 19.10.R6 for the 7450 Ethernet Service Switch (ESS), 7750 Service Router (SR), and 7950 eXtensible Routing System (XRS) platforms.

This document also provides an overview of the Virtualized Service Router (VSR). The VSR allows SR OS software to run in a virtual machine (VM) hosted on a standard x86 compute server. The VSR allows service providers to use Network Function Virtualization (NFV) to deliver the same types of services and functions available on 7750 SR routers. This release supports only VSR-I systems (use one virtual machine).

1.2 Release Numbering Model

Release 19.5.R1 introduced a new numbering model for 7450 ESS, 7750 SR, 7950 XRS, and VSR software releases. The new numbering model aligns these releases with the numbering of other Nokia products, and better communicates the timeframe in which the SR OS image was released.

In the new numbering convention:

- The major release number uses the last two digits of the year in which it is released. For example, the 2019 releases use the major release number 19.x.
- The minor release number uses the month number in which the feature release was initially made available. For example, a feature release made available in May 2019 has the release number 19.5, where 5 indicates the fifth month. The minor release number remains fixed even if the minor release is delayed past the end of the intended month.
- In addition, the SR OS version numbers have a maintenance release number. The initial R for a feature release is R1. Only the R1 feature releases will have new capabilities and enhancements, while subsequent maintenance releases of the same minor release will only contain bug fixes, and also small enhancements in R3 and higher.

As with previous releases of SR OS, In-Service Software Upgrade (ISSU) and associated features are supported with the final feature release in a specific year. In 2019, this was Release 19.10.R1.

1.3 Release 19.10.R6 Documentation Set

The SR OS Release 19.10.R6 documentation set consists of Release Notes and the 7450 ESS, 7750 SR, 7950 XRS, and VSR user guides. The components of the documentation set are listed in Table 1. New guides introduced in Release 19.*x* are highlighted in **bold**.

Table 1 **Release 19.10.R6 Documentation Set**

Document title	Part number
SR OS 19.10.R6 Software Release Notes	3HE 15407 0010 TQZZA
SR OS AA Applications and Protocols Release Notes R19.x	3HE 15408 0000 TQZZA
Acronyms Reference Guide	3HE 15078 AAAC TQZZA
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 19.10.R1 - Part I	3HE 14990 AAAC TQZZA
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 19.10.R1 - Part II	3HE 14991 AAAC TQZZA
Advanced Configuration Guide for 7450 ESS, 7750 SR and 7950 XRS for Releases up to 19.10.R1 - Part III	3HE 14992 AAAC TQZZA
vSIM Installation and Setup Guide	3HE 15073 AAAC TQZZA
VSR Installation and Setup Guide	3HE 15074 AAAC TQZZA
VSR Documentation Suite Overview	3HE 15075 AAAC TQZZA
Basic System Configuration Guide	3HE 15079 AAAC TQZZA
Documentation Suite Overview	3HE 15080 AAAC TQZZA
Gx AVPs Reference Guide	3HE 15081 AAAC TQZZA
Gy AVPs Reference Guide	3HE 15082 AAAC TQZZA
Interface Configuration Guide	3HE 15083 AAAC TQZZA
Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN	3HE 15084 AAAC TQZZA

Table 1 **Release 19.10.R6 Documentation Set (Continued)**

Document title	Part number
Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services	3HE 15085 AAAC TQZZA
Log Events Guide	3HE 15086 AAAC TQZZA
MD-CLI Command Reference Guide	3HE 15087 AAAJ TQZZA
MD-CLI Quick Reference Guide	3HE 15218 0000 TQZZA
MD-CLI User Guide	3HE 15088 AAAD TQZZA
MPLS Guide	3HE 15089 AAAC TQZZA
Multicast Routing Protocols Guide	3HE 15090 AAAC TQZZA
Multiservice Integrated Service Adapter and Extended Service Appliance Guide	3HE 15091 AAAC TQZZA
OAM and Diagnostics Guide	3HE 15092 AAAC TQZZA
Quality of Service Guide	3HE 15093 AAAC TQZZA
RADIUS Attributes Reference Guide	3HE 15094 AAAC TQZZA
Router Configuration Guide	3HE 15095 AAAC TQZZA
Services Overview Guide	3HE 15096 AAAC TQZZA
System Management Guide	3HE 15097 AAAC TQZZA
Triple Play Service Delivery Architecture Guide	3HE 15098 AAAC TQZZA
Unicast Routing Protocols Guide	3HE 15099 AAAC TQZZA
Zipped collection of documents	3HE 15100 AAAC TQZZA
Downloadable HTML Collection	3HE 15102 AAAC TQZZA
7450 ESS and 7750 SR Troubleshooting Guide	3HE 11475 AAAC TQZZA
Pay-As-You-Grow FP4 Hardware Licensing Reference Guide	**3HE 15453 AAAB TQZZA**

1.4 Guide Conventions

This guide uses the following terminology:

- SR OS node—the 7450 ESS, 7750 SR, 7950 XRS, and VSR platforms

- 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 16.0 and later. Nokia 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
- VSR may be used interchangeably with VSR-I, since only VSR-I is supported.

2 Release 19.10.R6 Supported Hardware

The following tables summarize the hardware supported in SR OS Release 19.10.R6. New hardware supported in SR OS Release 19.*x* is printed in **bold**.

➡ **Note:** QSFP28 optical modules are supported on some newer MDA-e-XP, XMA, and XMA-s assemblies. In addition, QSFP-DD optical modules are supported on some newer XMA and XMA-s assemblies. These optical modules can be single-port connectors or multi-port connectors. The number of 'ports' in the description of these assemblies refers to the case of only single port connectors being used. If multi-port connectors are used, then the number of physical Ethernet ports may be larger than what is listed in the description of the assembly. For example, the "XMA - SR-s 2.4T 36pt QSFP28 HE up to 3.6T" assembly could have as many as 360 Ethernet ports if a 10x10GE multi-port connector is used in every QSFP28 module cage on the card.

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
7450 ESS-7	7450 ESS-7 chassis (AC and DC)

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

Nokia Model #	Description
7450 ESS-12	7450 ESS-12 chassis (AC and DC)
7750 SR-1	7750 SR-1 chassis (AC and DC)
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-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)
7750 SR-1s	7750 SR-1s chassis (LVDC or AC/HVDC)
7750 SR-2s	7750 SR-2s chassis (LVDC or AC/HVDC)
7750 SR-7s	7750 SR-7s chassis (Externally mounted LVDC or AC/HVDC Power shelf)
7750 SR-14s	7750 SR-14s chassis (Externally mounted LVDC or AC/HVDC Power shelf)

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

The following tables summarize the Switch Fabric and Control Processor Modules (SFMs or CPMs), XMA Control Modules (XCMs), Connection and Control Modules (CCMs), CPM Module Adapter (CMA), 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	XRS-16c	XRS-20	XRS-20e	XRS-40	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
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
3HE12313AA	SFM - 7950 XRS-20/20e SFM2-SA		✓	✓		sfm2-x20s
3HE12321AA	XCM - 7950 XRS-20e XCM2			✓		xcm2-x20
3HE12327AA	XCM - 7950 XRS-20 XCM2		✓			xcm2-x20

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE03619AA	7750 SR IOM3-XP			✓	✓				✓	iom3-xp	--
3HE03620AA	IOM - 7450 ESS IOM3-XP								✓	iom3-xp	--
3HE03622AA	7750 SR 4pt 10GE XFP IMM			✓					✓	imm4-10gb-xfp	imm2-10gb-xp-xfp imm2-10gb-xp-xfp
3HE03623AA	7750 SR 8pt 10GE XFP IMM			✓					✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE03624AA	7750 SR 48pt GE SFP IMM			✓	✓				✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE03625AA	7750 SR 48pt GE copper/TX IMM			✓	✓				✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE04741AA	7750 SR 5pt 10GE XFP IMM			✓	✓				✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE04743AA	7x50 12pt 10G Ethernet SFP+ IMM – L3HQ			✓					✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05053AA	7x50 1pt 100G Ethernet CFP IMM – L3HQ			✓					✓	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05553AA	7x50 12pt 10G Ethernet SFP+ IMM – L2HQ			✓					✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05553BA	7x50 12pt 10G Ethernet SFP+ IMM – L3BQ			✓					✓	imm12-10gb-sf+	imm12-10gb-xp-sf+
3HE05814AA	7x50 1pt 100G Ethernet CFP IMM – L2HQ			✓					✓	imm1-100gb-cfp	imm1-100gb-xp-cfp

Table 5 SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE05814BA	7x50 1pt 100G Ethernet CFP IMM – L3BQ			✓					✓	imm1-100gb-cfp	imm1-100gb-xp-cfp
3HE05895AA	7x50 48pt GE SFP IMM – L2HQ			✓	✓				✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05895BA	7x50 48pt GE SFP IMM – L3BQ			✓	✓				✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE05896AA	7x50 48pt GE copper/ TX IMM – L2HQ			✓	✓				✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05896BA	7x50 48pt GE copper/ TX IMM – L3BQ			✓	✓				✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE05898AA	7x50 5pt 10GE XFP IMM – L2HQ			✓	✓				✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05898BA	7x50 5pt 10GE XFP IMM – L3BQ			✓	✓				✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE05899AA	7x50 8pt 10GE XFP IMM – L2HQ			✓					✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE05899BA	7x50 8pt 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	--

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE06320AA	7x50 3pt 40GE QSFP IMM – L3HQ			✓					✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE06324AA	IOM - 7450 ESS IOM3-XP (MULTI-CORE CPU)								✓	iom3-xp-b	--
3HE06326AA	7x50 48pt GE Multicore-CPU SFP IMM – L3HQ			✓	✓				✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326BA	7x50 48pt GE Multicore-CPU SFP IMM – L3BQ			✓	✓				✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06326CA	7x50 48pt GE Multicore-CPU SFP IMM – L2HQ			✓	✓				✓	imm48-1gb-sfp-b	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06428AA	7x50 48pt GE SFP IMM – L3HQ			✓	✓				✓	imm48-1gb-sfp	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE06429AA	7x50 48pt GE copper/ TX IMM – L3HQ			✓	✓				✓	imm48-1gb-tx	imm24-1gb-xp-tx imm24-1gb-xp-tx
3HE06430AA	7x50 5pt 10GE XFP IMM – L3HQ			✓	✓				✓	imm5-10gb-xfp	imm5-10gb-xp-xfp
3HE06431AA	7x50 8pt 10GE XFP IMM – L3HQ			✓					✓	imm8-10gb-xfp	imm4-10gb-xp-xfp imm4-10gb-xp-xfp
3HE06721AA	7x50 3pt 40GE QSFP IMM – L2HQ			✓					✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE06721BA	7x50 3pt 40GE QSFP IMM – L3BQ			✓					✓	imm3-40gb-qsfp	imm3-40gb-xp-qsfp
3HE07158AA	7x50 12pt 10GE FP3 SFP+ IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158BA	7x50 12pt 10GE FP3 SFP+ IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07158CA	7x50 12pt 10GE FP3 SFP+ IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p6-10g-sfp p6-10g-sfp
3HE07159AA	7x50 1pt 100GE FP3 CFP IMM – L3HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-cfp
3HE07159BA	7x50 1pt 100GE FP3 CFP IMM – L3BQ			✓	✓				✓	imm-1pac-fp3	p1-100g-cfp
3HE07159CA	7x50 1pt 100GE FP3 CFP IMM – L2HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-cfp
3HE07303AA	7x50 2pt 100GE FP3 CFP IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303BA	7x50 2pt 100GE FP3 CFP IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07303CA	7x50 2pt 100GE FP3 CFP IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p1-100g-cfp
3HE07304AA	7x50 6pt 40GE FP3 QSFP IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304BA	7x50 6pt 40GE FP3 QSFP IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07304CA	7x50 6pt 40GE FP3 QSFP IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p3-40g-qsfp
3HE07305AA	7x50 20pt 10GE FP3 SFP+ IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE07305BA	7x50 20pt 10GE FP3 SFP+ IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE07305CA	7x50 20pt 10GE FP3 SFP+ IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p10-10g-sfp
3HE08019AA	7x50 1pt 100GE DWDM Tunable FP3 IMM – L3HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun
3HE08019BA	7x50 1pt 100GE DWDM Tunable FP3 IMM – L3BQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun
3HE08019CA	7x50 1pt 100GE DWDM Tunable FP3 IMM – L2HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun
3HE08020AA	7x50 1pt 100GE CFP + 10pt 10GE SFP+ FP3 IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020BA	7x50 1pt 100GE CFP + 10pt 10GE SFP+ FP3 IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08020CA	7x50 1pt 100GE CFP + 10pt 10GE SFP+ FP3 IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p1-100g-cfp p10-10g-sfp
3HE08174AA	7x50 10pt 10GE SFP+ + 20pt GE SFP FP3 IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174BA	7x50 10pt 10GE SFP+ + 20pt GE SFP FP3 IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp
3HE08174CA	7x50 10pt 10GE SFP+ + 20pt GE SFP FP3 IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p10-10g-sfp p20-1ge-sfp

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE08175AA	7x50 3pt 40GE QSFP + 20pt GE SFP FP3 IMM – L3HQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175BA	7x50 3pt 40GE QSFP + 20pt GE SFP FP3 IMM – L3BQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08175CA	7x50 3pt 40GE QSFP + 20pt GE SFP FP3 IMM – L2HQ			✓	✓				✓	imm-2pac-fp3	p3-40g-qsfp p20-1ge-sfp
3HE08421AA	7750 SR SFM5-12e				✓					m-sfm5-12e	--
3HE08422AA	7750 SR Mini-SFM5-12e				✓					m-sfm5-12e	--
3HE08423AA	7750 SR CPM5			✓	✓					cpm5	--
3HE08424AA	7x50 40pt 10GE SFP+ IMM – L3HQ				✓					imm40-10gb-sfp	m40-10g-sfp
3HE08424BA	7x50 40pt 10GE SFP+ IMM – L3BQ				✓					imm40-10gb-sfp	m40-10g-sfp
3HE08424CA	7x50 40pt 10GE SFP+ IMM – L2HQ				✓					imm40-10gb-sfp	m40-10g-sfp
3HE08425AA	7x50 4pt 100GE CXP IMM – L3HQ				✓					imm4-100gb-cxp	m4-100g-cxp
3HE08425BA	7x50 4pt 100GE CXP IMM – L3BQ				✓					imm4-100gb-cxp	m4-100g-cxp
3HE08425CA	7x50 4pt 100GE CXP IMM – L2HQ				✓					imm4-100gb-cxp	m4-100g-cxp
3HE08426AA	7750 SR IOM3-XP-C			✓	✓				✓	iom3-xp-c	--
3HE08427AA	7450 ESS IOM3-XP-C								✓	iom3-xp-c	--
3HE08428AA	7750 SR SFM5-12 [1]			✓						m-sfm5-12	--
3HE08429AA	7750 SR SFM5-7 [2]			✓						m-sfm5-7	--

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE08430AA	7450 ESS SFM5-12 [3]								✓	sfm5-12	--
3HE08431AA	7450 ESS SFM5-7 [4]								✓	sfm5-7	--
3HE08432AA	7450 ESS CPM5								✓	cpm5	--
3HE09117AA	7x50 Multiservice ISM [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p-isa2-ms
3HE09118AA	7x50 Multiservice ISME (no encryption) [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p-isa2-ms-e
3HE09192AA	7x50 MS-ISA2 + 1pt 100GE CFP IMM – L3HQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192BA	7x50 MS-ISA2 + 1pt 100GE CFP IMM – L3BQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09192CA	7x50 MS-ISA2 + 1pt 100GE CFP IMM – L2HQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p1-100g-cfp
3HE09193AA	7x50 MS-ISA2 + 10pt 10GE SFP+ IMM – L3HQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193BA	7x50 MS-ISA2 + 10pt 10GE SFP+ IMM – L3BQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09193CA	7x50 MS-ISA2 + 10pt 10GE SFP+ IMM – L2HQ [7]			✓	✓				✓	imm-2pac-fp3	p-isa2-ms p10-10g-sfp
3HE09201AA	7750 SR-a CPM	✓								cpm-a	--
3HE09201AB	**7750 SR-a CPM 8GB**	✓								**cpm-a**	--
3HE09202AA	7750 SR-a IOM – L3HQ	✓								iom-a	--

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE09202BA	7750 SR-a IOM – L3BQ	✓								iom-a	--
3HE09202CA	7750 SR-a IOM – L2HQ	✓								iom-a	--
3HE09253AA	7x50 MS-ISA2-E + 1pt 100GE CFP IMM – L3HQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253BA	7x50 MS-ISA2-E + 1pt 100GE CFP IMM – L3BQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09253CA	7x50 MS-ISA2-E + 1pt 100GE CFP IMM – L2HQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p1-100g-cfp
3HE09254AA	7x50 MS-ISA2-E + 10pt 10G SFP+ IMM – L3HQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254BA	7x50 MS-ISA2-E + 10pt 10G SFP+ IMM – L3BQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09254CA	7x50 MS-ISA2-E + 10pt 10G SFP+ IMM – L2HQ [7]			✓	✓					imm-2pac-fp3	p-isa2-ms-e p10-10g-sfp
3HE09279AA	7x50 48pt GE Multicore SFP IMM – L3HQ			✓	✓				✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279BA	7x50 48pt GE Multicore SFP IMM – L3BQ			✓	✓				✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp
3HE09279CA	7x50 48pt GE Multicore SFP IMM – L2HQ			✓	✓				✓	imm48-1gb-sfp-c	imm24-1gb-xp-sfp imm24-1gb-xp-sfp

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE09436AA	IMM – 7750 SR 1pt 100GE INT DWDM L3HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436BA	IMM – 7750 SR 1pt 100GE INT DWDM L3BQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun-b
3HE09436CA	IMM – 7750 SR 1pt 100GE INT DWDM L2HQ			✓	✓				✓	imm-1pac-fp3	p1-100g-tun-b
3HE09645AA	7x50 4pt 100GE CFP4 IMM – L3HQ				✓					imm4-100gb-cfp4	m4-100g-cfp4
3HE09645BA	7x50 4pt 100GE CFP4 IMM – L3BQ				✓					imm4-100gb-cfp4	m4-100g-cfp4
3HE09645CA	7x50 4pt 100GE CFP4 IMM – L2HQ				✓					imm4-100gb-cfp4	m4-100g-cfp4
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 – 160pt GE cSFP/ 80pt GE SFP – L3HQ			✓	✓				✓	imm-1pac-fp3	p160-1gb-csfp
3HE10014BA	IMM – 160pt GE cSFP/ 80pt GE SFP – L3BQ			✓	✓				✓	imm-1pac-fp3	p160-1gb-csfp
3HE10014CA	IMM – 160pt GE cSFP/ 80pt 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	--

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
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	--
3HE11304AA	CPM - 7750 SR-s CPM-14s/7s						✓	✓		cpm-s	--
3HE11305AA	XCM - 7750 SR-s XCM-7s						✓			xcm-7s	--
3HE11315AA	SFM - 7750 SR-14s/7s						✓	✓		sfm-s	--
3HE11316AA	XCM - 7750 SR-s XCM-14s							✓		xcm-14s	--
3HE11351AA	IOM - 7750 SR IOM4-e-HS L3HQ			✓	✓					iom4-e-hs	--
3HE11351CA	IOM - 7750 SR IOM4-e-HS L2HQ			✓	✓					iom4-e-hs	--
3HE12314AA	CMA - 7750 SR-7s, dual CMA						✓			--	--
3HE12315AA	CMA - 7750 SR-14s							✓		--	--
3HE12326AA	CMA - 7750 SR-7s, single CMA						✓			--	--
3HE12330AA	SFM - 7750 SR SFM6-12e				✓					m-sfm6-12e	--

Table 5 **SFM, CPM, CCM, IOM, IMM, ISM, CMA, and XCM Cards Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	SR-a4/a8	SR-1e/2e/3e	SR-7/12	SR-12e	SR-2s	SR-7s	SR-14s	ESS-7/12	CLI String (Card)	CLI String (MDA)
3HE12331AA	SFM - 7750 SR Mini-SFM6-12e				✓					m-sfm6-12e	--
3HE12332AA	7750 SR IOM5-e – CR [5]			✓	✓					iom5-e	--
3HE12332BA	7750 SR IOM5-e – ER [5]			✓	✓					iom5-e	--
3HE12332CA	7750 SR IOM5-e – HE [5]			✓	✓					iom5-e	--
3HE12335AA	SFM - 7750 SR SFM6-7/12 [6]			✓						m-sfm6-7/12	--
3HE12379AA	XCM - 7750 SR-s XCM-2s					✓				xcm-2s	--
3HE13727AA	IOM - 7750 SR IOM5-e 400G - CR [5]			✓	✓					iom5-e	--
3HE13727BA	IOM - 7750 SR IOM5-e 400G - ER [5]			✓	✓					iom5-e	--
3HE13727CA	IOM - 7750 SR IOM5-e 400G – HE [5]			✓	✓					iom5-e	--
3HE14034AA	CPM - 7750 SR-2s CPM-2s					✓				cpm-2s	--
3HE15182AA	IOM - 7750 SR IOM5-e 800G - CR [5]			✓	✓					iom5-e	--
3HE15182BA	IOM - 7750 SR IOM5-e 800G - ER [5]			✓	✓					iom5-e	--
3HE15182CA	IOM - 7750 SR IOM5-e 800G – HE [5]			✓	✓					iom5-e	--

Notes:

1. Not supported on 7750 SR-7.
2. Not supported on 7750 SR-12.

3. Not supported on 7450 ESS-7.
4. Not supported on 7450 ESS-12.
5. IOM5-e cards for 7750 SR-7/12 are only supported on SR-7-B and SR-12-B.
6. Only supported on 7750 SR-7-B and SR-12-B.
7. See Usage Notes for specific details.

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

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

Table 6 **XMAs and C-XMAs Supported in 7950 XRS**

Nokia Part #	Description	XRS-16c	XRS-20	XRS-20e	XRS-40	CLI String (MDA)
3HE06937AA	C-XMA – 7950 XRS 20pt 10GE SFP+ – IP Core	✓	✓[1]	✓[1]	✓[1]	cx20-10g-sfp
3HE06938AA	C-XMA – 7950 XRS 2pt 100GE CFP – IP Core	✓	✓[1]	✓[1]	✓[1]	cx2-100g-cfp
3HE06937BA	C-XMA – 7950 XRS 20pt 10GE SFP+ – LSR	✓	✓[1]	✓[1]	✓[1]	cx20-10g-sfp
3HE06938BA	C-XMA – 7950 XRS 2pt 100GE CFP – LSR	✓	✓[1]	✓[1]	✓[1]	cx2-100g-cfp
3HE07297AA	XMA – 7950 XRS 40pt 10GE SFP+ – IP Core		✓	✓	✓	x40-10g-sfp
3HE07297BA	XMA – 7950 XRS 40pt 10GE SFP+ – LSR		✓	✓	✓	x40-10g-sfp
3HE07299AA	XMA – 7950 XRS 4pt 100GE CXP – IP Core		✓	✓	✓	x4-100g-cxp
3HE07299BA	XMA – 7950 XRS 4pt 100GE CXP – LSR		✓	✓	✓	x4-100g-cxp
3HE08214AA	C-XMA – 7950 XRS 6pt 40GE QSFP+ – IP Core	✓	✓[1]	✓[1]	✓[1]	cx6-40g-qsfp
3HE08214BA	C-XMA – 7950 XRS 6pt 40GE QSFP+ – LSR	✓	✓[1]	✓[1]	✓[1]	cx6-40g-qsfp
3HE08631AA	C-XMA – 7950 XRS 72pt GE CSFP – IP Core	✓	✓[1]	✓[1]	✓[1]	cx72-1g-csfp
3HE08631BA	C-XMA – 7950 XRS 72pt GE CSFP – LSR	✓	✓[1]	✓[1]	✓[1]	cx72-1g-csfp
3HE08632AA	XMA – 7950 XRS 4pt 100G CFP2 – IP Core		✓	✓	✓	x4-100g-cfp2
3HE08632BA	XMA – 7950 XRS 4pt 100G CFP2 – LSR		✓	✓	✓	x4-100g-cfp2

Table 6 **XMAs and C-XMAs Supported in 7950 XRS (Continued)**

Nokia Part #	Description	XRS-16c	XRS-20	XRS-20e	XRS-40	CLI String (MDA)
3HE10677AA	XMA – 7950 XRS 2pt 100GE INT DWDM – IP Core		✓	✓	✓	x2-100g-tun
3HE10677AB	XMA – 7950 XRS 2pt 100GE INT DWDM – LSR		✓	✓	✓	x2-100g-tun
3HE11094AA	XMA – 7950 XRS 2.4T 24pt QSFP28 – CR		✓[2]	✓		x24-100g-qsfp28
3HE11094BA	XMA – 7950 XRS 2.4T 24pt QSFP28 – ER		✓[2]	✓		x24-100g-qsfp28
3HE11094CA	XMA – 7950 XRS 2.4T 24pt QSFP28 – HE		✓[2]	✓		x24-100g-qsfp28
3HE12324AA	XMA - 7950 XRS 6pt CFP8 XMA – CR		✓[2]	✓		x6-400g-cfp8
3HE12324BA	XMA - 7950 XRS 6pt CFP8 XMA – ER		✓[2]	✓		x6-400g-cfp8
3HE12324CA	XMA - 7950 XRS 6pt CFP8 XMA – HE		✓[2]	✓		x6-400g-cfp8
3HE13810AA	XMA – 7950 XRS 1.2T 12pt QSFP28 – CR		✓	✓		x24-100g-qsfp28
3HE13810BA	XMA – 7950 XRS 1.2T 12pt QSFP28 – ER		✓	✓		x24-100g-qsfp28
3HE13810CA	XMA – 7950 XRS 1.2T 12pt QSFP28 – HE		✓	✓		x24-100g-qsfp28
3HE13811AA	XMA – 7950 XRS 1.6T 16pt QSFP28 – CR		✓	✓		x24-100g-qsfp28
3HE13811BA	XMA – 7950 XRS 1.6T 16pt QSFP28 – ER		✓	✓		x24-100g-qsfp28
3HE13811CA	XMA – 7950 XRS 1.6T 16pt QSFP28 – HE		✓	✓		x24-100g-qsfp28
3HE13812AA	XMA - 7950 XRS 4pt CFP8 XMA – CR		✓	✓		x6-400g-cfp8
3HE13812BA	XMA - 7950 XRS 4pt CFP8 XMA – ER		✓	✓		x6-400g-cfp8
3HE13812CA	XMA - 7950 XRS 4pt CFP8 XMA – HE		✓	✓		x6-400g-cfp8
3HE14662AA	**XMA - 7950 XRS 1.6T 8pt QSFP-DD – CR**		✓	✓		**x12-400g-qsfpdd**
3HE14662BA	**XMA - 7950 XRS 1.6T 8pt QSFP-DD – ER**		✓	✓		**x12-400g-qsfpdd**
3HE14662CA	**XMA - 7950 XRS 1.6T 8pt QSFP-DD – HE**		✓	✓		**x12-400g-qsfpdd**
3HE14663AA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD – CR**		✓[2]	✓		**x12-400g-qsfpdd**
3HE14663BA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD – ER**		✓[2]	✓		**x12-400g-qsfpdd**
3HE14663CA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD – HE**		✓[2]	✓		**x12-400g-qsfpdd**

Table 6 XMAs and C-XMAs Supported in 7950 XRS (Continued)

Nokia Part #	Description	XRS-16c	XRS-20	XRS-20e	XRS-40	CLI String (MDA)
3HE14664AA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD to 4T – CR**		✓[2]	✓		**x12-400g-qsfpdd**
3HE14664BA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD to 4T – ER**		✓[2]	✓		**x12-400g-qsfpdd**
3HE14664CA	**XMA - 7950 XRS 2.4T 12pt QSFP-DD to 4T – HE**		✓[2]	✓		**x12-400g-qsfpdd**
3HE14665AA	**XMA - 7950 XRS 6pt CFP2-DCO – CR**		✓[3]	✓		**x6-200g-cfp2-dco**
3HE14665BA	**XMA - 7950 XRS 6pt CFP2-DCO – ER**		✓[3]	✓		**x6-200g-cfp2-dco**
3HE14665CA	**XMA - 7950 XRS 6pt CFP2-DCO – HE**		✓[3]	✓		**x6-200g-cfp2-dco**

Notes:

1. C-XMA adapter is required.
2. Connector and bandwidth limitations apply when installed in an XRS-20.
3. Bandwidth limitations apply when installed in an XRS-20.

Table 7 MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE00021AA	60pt 10/100TX MDA – mini-RJ21	✓					m60-10/100eth-tx
3HE00230AA	7450 ESS 60pt 10/100TX MDA [5]	✓					m60-10/100eth-tx
3HE01364AA	4pt Channelized OC-3/STM-1 (DS0) ASAP MDA – SFP	✓					m4-choc3-as-sfp
3HE02021AA	1pt 10GBASE + 10pt GIGE MDA	✓					m10-1gb+1-10gb

Table 7 **MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE02022AA	7450 ESS 1pt 10GE+10P GIGE-NO OPT [5]	✓					m10-1gb+1-10gb
3HE02499AA	1pt Channelized OC-12/STM-4 ASAP MDA	✓					m1-choc12-as-sfp
3HE02500AA	12pt Channelized DS3/E3 ASAP MDA	✓					m12-chds3-as
3HE02501AA	4pt Channelized DS3/E3 ASAP MDA	✓					m4-chds3-as
3HE03078AA	1pt Channelized OC-3/STM-1 CES MDA	✓					m1-choc3-ces-sfp
3HE03079AA	7750 SR 4pt CH OC-3/STM-1 CES SFP MDA	✓					m4-choc3-ces-sfp
3HE03611AA	7750 SR 10pt GE – SFP MDA-XP	✓					m10-1gb-xp-sfp
3HE03612AA	7750 SR 20pt GE – SFP MDA-XP	✓					m20-1gb-xp-sfp
3HE03613AA	7750 SR 20pt GE – Copper/TX MDA-XP	✓					m20-1gb-xp-tx
3HE03614AA	MDA - 7450 10pt GE MDA - XP SFP [5]	✓					m10-1gb-xp-sfp
3HE03615AA	7450 ESS 20-port GE - SFP MDA-XP	✓					m20-1gb-xp-sfp
3HE03616AA	MDA - 7450 20pt GE – Copper/TX MDA-XP	✓					m20-1gb-xp-tx
3HE03685AA	7750 SR 2pt 10GBASE – XFP MDA-XP	✓					m2-10gb-xp-xfp
3HE03686AA	7750 SR 4pt 10GBASE – XFP MDA-XP	✓					m4-10gb-xp-xfp
3HE03687AA	MDA - 7450 ESS 2pt 10G MDA-XP - XFP [5]	✓					m2-10gb-xp-xfp
3HE03688AA	MDA - 7450 ESS 4pt 10G MDA-XP - XFP [5]	✓					m4-10gb-xp-xfp
3HE04272AA	7750 SR 1pt OC-12/STM-4 CES MDA	✓					m1-choc12-ces-sfp
3HE04273AA	MDA - 7450 ESS 1pt 10G MDA-XP - XFP [5]	✓					m1-10gb-xp-xfp
3HE04274AA	7750 SR 1pt 10GBASE – XFP MDA-XP	✓					m1-10gb-xp-xfp
3HE04922AA	7750 SR / 7450 ESS Multiservice ISA [1]	✓					isa-ms

Table 7 **MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE05142AA	7750 SR / 7450 ESS Multiservice ISA-E (no encryption)[1]	✓					isa-ms-e
3HE05159AA	MDA-XP 7450 ESS 48pt 10/100/1000 MRJ-21[5]	✓					m48-1gb-xp-tx
3HE05160AA	7750 SR 48pt 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 16pt OC-3/12c STM-1/4c POS MDA – SFP Rev B	✓					m16-oc12/3-sfp-b
3HE05946AA	7750 SR 4pt OC-48c/STM-16c POS MDA – SFP Rev B	✓					m4-oc48-sfp-b
3HE05947AA	7750 SR 2pt OC-192/STM-64 – XFP MDA-XP	✓					m2-oc192-xp-xfp
3HE06432AA	7750 SR 10pt GE SFP HS-MDAv2	✓					m10-1gb-hs-sfp-b
3HE06433AA	7750 SR 1pt 10GE HS-MDAv2	✓					m1-10gb-hs-xfp-b
3HE06434AA	MDA - 7450 ESS 10pt GIGE HS-MDA2[5]	✓					m10-1gb-hs-sfp-b
3HE06435AA	7450 ESS 1pt 10GE HS-MDAv2[5]	✓					m1-10gb-hs-xfp-b
3HE07282AA	7750 SR 2pt 10GE XFP + 12pt GE SFP – MDA-XP	✓					m12-1gb+2-10gb-xp
3HE07283AA	MDA 7450 2x10GE XFP + 12x1GE SFP MDA-XP[5]	✓					m12-1gb+2-10gb-xp
3HE07284AA	7750 SR 12pt GigE – SFP MDA-XP	✓					m12-1gb-xp-sfp
3HE07285AA	MDA - 7450 ESS 12pt 1GE MDA-XP[5]	✓					m12-1gb-sfp
3HE09203AA	7750 SR-a 1pt 100GE MDA-a XP – CFP			✓			maxp1-100gb-cfp

Table 7 **MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE09204AA	7750 SR-a 10pt 10GE MDA-a XP – SFP+			✓			maxp10-10gb-sfp+
3HE09205AA	7750 SR-a 2pt 10GE SFP+ + 12pt GE SFP MDA-a			✓			ma2-10gb-sfp+12-1gb-sfp
3HE09206AA	7750 SR-a 20pt 10/100/1000 TX MDA-a – RJ45			✓			ma20-1gb-tx
3HE09207AA	7750 SR-a 22pt GE SFP/44pt GE MDA-a – CSFP			✓			ma44-1gb-csfp
3HE09240AA	7750 SR-a 4pt 10GE MDA-a – SFP+			✓			ma4-10gb-sfp+
3HE09241AA	7750 SR-a 6pt 10GE SFP+ + 1pt 40GE QSFP+ MDA-a XP			✓			maxp6-10gb-sfp+1-40gb-qsfp+
3HE09649AA	MDA-e 10pt 10 GE SFP+		✓				me10-10gb-sfp+
3HE09881AA	MDA-e 1pt 100 GE CFP2		✓				me1-100gb-cfp2
3HE10421AA	MDA–a XP - 7750 SR 1pt 100GE CFP2			✓			maxp1-100gb-cfp2
3HE10422AA	MDA–a XP - 7750 SR 1pt 100GE CFP4			✓			maxp1-100gb-cfp4
3HE10427AA	ISA - 7750 SR MS-ISA2 [1,2]		✓				me-isa2-ms
3HE10428AA	ISA - 7750 SR MS-ISA2-E [1,2]		✓				me-isa2-ms-e
3HE10429AA	MDA-e 6pt 10GE SFP+		✓				me6-10gb-sfp+
3HE10642AA	MDA-e 20pt GE SFP/40pt GE cSFP		✓				me40-1gb-csfp
3HE11030AA	MDA-e 2pt 100GE CFP4		✓				me2-100gb-cfp4
3HE11031AA	MDA-e 2pt 100GE QSFP28		✓				me2-100gb-qsfp28
3HE11307AA	XMA - SR-s 2.4T 36pt QSFP28 CR up to 3.6T					✓	s36-100gb-qsfp28
3HE11307BA	XMA - SR-s 2.4T 36pt QSFP28 ER up to 3.6T					✓	s36-100gb-qsfp28

Table 7 **MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE11307CA	XMA - SR-s 2.4T 36pt QSFP28 HE up to 3.6T					✓	s36-100gb-qsfp28
3HE11903AA	MDA-e 12pt 10/1GE MACsec SFP+		✓				me12-10/1gb-sfp+
3HE12333AA	MDA-e-XP - 7750 SR 6pt 100GE QSFP28				✓		me6-100gb-qsfp28
3HE12334AA	MDA-e-XP - 7750 SR 12pt 100GE QSFP28 [3]				✓		me12-100gb-qsfp28
3HE12388AA	XMA - SR-s 2.4T 24pt QSFP28 – CR					✓	s36-100gb-qsfp28
3HE12388BA	XMA - SR-s 2.4T 24pt QSFP28 – ER					✓	s36-100gb-qsfp28
3HE12388CA	XMA - SR-s 2.4T 24pt QSFP28 – HE					✓	s36-100gb-qsfp28
3HE12389AA	XMA - SR-s 3.6T 36pt QSFP28 – CR					✓	s36-400gb-qsfpdd
3HE12389BA	XMA - SR-s 3.6T 36pt QSFP28 – ER					✓	s36-400gb-qsfpdd
3HE12389CA	XMA - SR-s 3.6T 36pt QSFP28 – HE					✓	s36-400gb-qsfpdd
3HE12390AA	XMA - SR-s 3.6T 36pt QSFP-DD – CR					✓	s36-400gb-qsfpdd
3HE12390BA	XMA - SR-s 3.6T 36pt QSFP-DD – ER					✓	s36-400gb-qsfpdd
3HE12390CA	XMA - SR-s 3.6T 36pt QSFP-DD – HE					✓	s36-400gb-qsfpdd
3HE12391AA	XMA - SR-s 4.8T 36pt QSFP-DD – CR					✓	s36-400gb-qsfpdd
3HE12391BA	XMA - SR-s 4.8T 36pt QSFP-DD – ER					✓	s36-400gb-qsfpdd
3HE12391CA	XMA - SR-s 4.8T 36pt QSFP-DD – HE					✓	s36-400gb-qsfpdd
3HE12392AA	XMA - SR-s 4.8T 36pt QSFP-DD CR up to 12T					✓	s36-400gb-qsfpdd
3HE12392BA	XMA - SR-s 4.8T 36pt QSFP-DD ER up to 12T					✓	s36-400gb-qsfpdd
3HE12392CA	XMA - SR-s 4.8T 36pt QSFP-DD HE up to 12T					✓	s36-400gb-qsfpdd

Table 7 **MDAs, XMAs, and ISAs Supported in 7750 SR and 7450 ESS (Continued)**

Nokia Part #	Description	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, iom5-e	xcm-2s/7s/14s	CLI String (MDA)
3HE12407AA	**MDA-e - 7750 SR 2pt 100GE – Multi-Service QSFP28**		✓				**me2-100gb-ms-qsfp28**
3HE12524AA	XMA - SR-s 600G 6pt QSFP28 – CR					✓	s18-100gb-qsfp28
3HE12524BA	XMA - SR-s 600G 6pt QSFP28 – ER					✓	s18-100gb-qsfp28
3HE12524CA	XMA - SR-s 600G 6pt QSFP28 – HE					✓	s18-100gb-qsfp28
3HE12525AA	XMA - SR-s 1.2T 12pt QSFP28 – CR					✓	s18-100gb-qsfp28
3HE12525BA	XMA - SR-s 1.2T 12pt QSFP28 – ER					✓	s18-100gb-qsfp28
3HE12525CA	XMA - SR-s 1.2T 12pt QSFP28 – HE					✓	s18-100gb-qsfp28
3HE12526AA	XMA - SR-s 1.2T 18pt QSFP28 CR to 1.8T					✓	s18-100gb-qsfp28
3HE12526BA	XMA - SR-s 1.2T 18pt QSFP28 ER to 1.8T					✓	s18-100gb-qsfp28
3HE12526CA	XMA - SR-s 1.2T 18pt QSFP28 HE to 1.8T					✓	s18-100gb-qsfp28
3HE13740AA	XMA - SR-s 1.6T 16pt QSFP28 – CR					✓	s36-100gb-qsfp28
3HE13740BA	XMA - SR-s 1.6T 16pt QSFP28 – ER					✓	s36-100gb-qsfp28
3HE13740CA	XMA - SR-s 1.6T 16pt QSFP28 – HE					✓	s36-100gb-qsfp28
3HE13955AA	**MDA-e-XP - 7750 SR 3pt CFP2-DCO**				✓		**me3-200gb-cfp2-dco**
3HE14035AA	**MDA-a XP - 7750 SR 10pt 10/1GE MACsec**			✓			**maxp10-10/1gb-msec-sfp+**
3HE15152AA	**MDA-e-XP - 7750 SR 6pt QSFP-DD**				✓		**me6-400gb-qsfpdd**

Notes:

1. See Usage Notes for specifics.
2. Only Ethernet MDA-e cards are supported in IOM4-e-hs. ISAs are not supported.

3. Not supported on the 7750 SR-7-B/12-B chassis.
4. Not supported on the 7750 SR-1.
5. Not supported on the 7750 SR-7/12/12e.

2.4 Supported 7210 SAS-Sx Satellites

Table 8 summarizes the 7210 SAS-Sx satellites supported in SR OS.

Host SR OS Release 19.*x* supports 7210 Ethernet satellite images 9.0, 10.0, and 11.0.

Host SR OS Release 19.*x* supports 7705 SAR satellite images 8.0 and 9.0 on the SONET/SDH satellite. For Release 8.0, it must be 8.0.R4 or higher.

Table 8 **7210 SAS-Sx Satellites**

Nokia Part #	Description	sat-type	Initial 7210 Software Release	Initial 7750 Host Support
3HE10075AB	7210 SAS-Mxp 22F2C4SFP+	es24-sasmxp-1gb-sfp	10.0.R5	16.0.R2
3HE10076AB	7210 SAS-Mxp 22F2C4SFP+ ETR	es24-sasmxp-1gb-sfp	10.0.R5	16.0.R2
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
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

Table 8 **7210 SAS-Sx Satellites (Continued)**

Nokia Part #	Description	sat-type	Initial 7210 Software Release	Initial 7750 Host Support
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
3HE10538AB	SYS - 7210 SAS-S 48Tp4SFP+ (AC)	es48-1gb-tx	10.0.R8	16.0.R4
3HE10540AB	SYS - 7210 SAS-S 24Tp4SFP+ (AC)	es24-1gb-tx	10.0.R8	16.0.R4
3HE10835AA	SYS - 7210 SAS-Sx 64SFP+ 4CFP4	es64-10gb-sfpp+4-100gb-cfp4	9.0.R7	15.0.R4
3HE11597AA	SYS - 7210 SAS-Sx 64SFP+ 4QSFP28	es64-10gb-sfpp+4-100gb-qsfp28	10.0.R8	16.0.R4

3 New Features

The following sections describe the new features added in SR OS releases. New features from Releases 16.0.R2 to 16.0.R7 also apply to Release 19.*x*. Refer to the most recent *SR OS 16.0 Release Notes* for the summary of new features in Releases 16.0.R2 through 16.0.R7.

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 19.10.R6

There are no new major features in Release 19.10.R6. See also the Resolved Issues in Release 19.10.R6.

3.2 Release 19.10.R5

There are no new major features in Release 19.10.R5. See also the Resolved Issues in Release 19.10.R5.

3.3 Release 19.10.R4

There are no new major features in Release 19.10.R4. See also the Resolved Issues in Release 19.10.R4.

3.4 Release 19.10.R3

3.4.1 Hardware

3.4.1.1 7750 SR 6-port QSFP-DD MDA-e-XP

Release 19.10.R3 introduces support for the 7750 SR 6-port QSFP-DD MDA-e-XP on the 7750 SR-1/12e with the IOM5-e.

3.4.2 System

3.4.2.1 SSHv2 Configurable Key Exchange Lists

Release 19.10.R3 introduces the SSHv2 configurable Diffie-Hellman key exchange (KEX) lists for client and server. The configurable list adds or removes any unwanted algorithms from SSHv2 phase one negotiations. The list is an index list with the lower index having a higher priority in the SSH negotiation. In addition, SHA2 algorithms with DHG 2k and higher prime numbers were added to SSHv2 phase one KEX negotiation. These algorithms include diffie-hellman-group14-sha256 and diffie-hellman-group16-sha512. By default, these new algorithms are on top of the negotiation list, so they have the lowest index in the KEX server/client list.

3.4.3 MPLS

3.4.3.1 Traffic Statistics for Two Labels in the MPLS Stack

Release 19.10.R3 introduces the capability to control how many labels in the MPLS stack the system is able to collect traffic statistics on. Prior to release 19.10.R3, it was possible to collect traffic statistics on only one tunnel, even if multiple tunnel types were stacked and traffic statistics were enabled on several on them. Traffic statistics were collected only for the outermost label with statistics enabled. Release 19.10.R3 allows traffic statistics collection on one (default) or two labels in the MPLS stack. When traffic statistics on two labels is configured, the system will collect statistics on the outermost label with statistics enabled (backward compatible behavior) and the innermost label with statistics enabled.

3.4.3.2 Configure SID for Path Hops

Release 19.10.R3 adds the ability to use LSP Ping and LSP Trace on SR-TE LSP paths using explicit SID labels.

3.4.3.3 LDP ECMP Next-hop Limit Increased to 64 (Multicast)

Release 19.10.R3 increases the maximum number of ECMP next-hops available to LDP from 32 to 64. This is applicable to cases where LDP resolves a FEC over network IP interfaces, IGP shortcut tunnels with RSVP-TE and SR-TE LSPs, LDP-over-RSVP, and FEC 129 and P2MP FECs.

The **config**>**router**>**ldp**>**max-ecmp-routes** command is introduced to configure the increased limit, subject to a maximum set in the **config**>**router**>**ecmp** context. The default limit for LDP remains at 32.

In Release 19.10.R3, this increase is only introduced for Multicast. The increase was introduced for Unicast in Release 19.10.R1.

3.5 Release 19.10.R2

3.5.1 System

3.5.1.1 ISSU Across Minor Releases

In-Service Software Update (ISSU) across minor releases (Minor ISSU) allows in-service software updates across maintenance releases (within the same major release) for systems with dual CPMs without requiring a reboot of the system. ISSU is comparable to performing a controlled High-Availability switchover where the new image is loaded onto the standby CPM which becomes master, and then upgrading the image on the other CPM. The terms Major ISSU and Minor ISSU are used to differentiate between ISSU across major releases and maintenance releases within a major release, respectively.

See the ISSU Upgrade Procedure for more information on ISSU operation, including applicable upgrade paths.

3.5.1.2 Minor ISSU Support in Model-Driven Mode

Release 19.10.R2 introduces the support for Minor ISSU when **configure system management-interface configuration-mode** is set to **model-driven**. This supports handling of changed or removed commands between the releases involved in the ISSU as described in the Changed or Removed Commands section of the relevant Release Notes.

See the ISSU Upgrade Procedure for more information on ISSU operation, including applicable upgrade paths.

3.6 Release 19.10.R1

3.6.1 Hardware

3.6.1.1 3-port CFP2-DCO MDA-e-XP support on 7750 SR-7-B/12-B

Release 19.10.R1 introduces the support for the 7750 SR 3-port CFP2-DCO MDA-e-XP on the 7750 SR-7-B/12-B chassis with the IOM5-e. This MDA-e-XP is already supported in the 7750 SR-1 and 7750 SR-12e without any restrictions. The following configuration options are supported with the 3-port CFP2-DCO MDA-e-XP on the SR-7/12-B chassis specifically:

- Up to 2x MDA-e-XP 3-port CFP2-DCOs per IOM5-e
- With two of these MDA-e-XPs in place, optics options are:
  - 3x CFP2 (@100G each) in potentially all 6 ports, or
  - 1x CFP2-DCO in the middle position on each of the two MDA-e-XPs for a total of 2x 200G ports

System operation is restricted to a maximum of 40°C when this MDA-e-XP is in place.

3.6.1.2 7750 SR-a CPM 8GB

Release 19.10.R1 introduces the support for the 7750 SR-a CPM-a with 8GB of RAM for the 7750 SR-a4 and SR-a8 chassis.

3.6.1.3 QSFP28-to-SFP+/SFP28 Adapter

Release 19.10.R1 introduces the support for the QSFP28-to-SFP+/SFP28 adapter for the 7750 SR platform on me6-100gb-qsfp28, me12-100gb-qsfp28, and s36-400gb-qsfpdd MDA types. The QSFP28-to-SFP+/SFP28 adapter provides 10GE (SFP+) support in a QSFP28/QSFP-DD connector. When a QSFP28 or QSFP-DD connector uses an SFP+ optical module with the QSFP28-to-SFP+/SFP28 adapter, the breakout should be set to **c1-10g**.

3.6.1.4 c1-400g (400GE support)

Release 19.10.R1 introduces **c1-400g** (400GE support) support for the x12-400g-qsfpdd (7950 XRS) and the s36-400gb-qsfpdd (7750 SR-s).

3.6.2 System

3.6.2.1 In-Service Software Upgrade across Major Releases (Major ISSU)

In-Service Software Upgrade (ISSU) allows in-service updates across one or two major releases for systems with dual-CPMs without requiring a reboot of the system. ISSU is comparable to performing a controlled High-Availability switchover where the new image is loaded onto the standby CPM which becomes master, and then upgrading the image on the other CPM. See the ISSU Upgrade Procedure, for more information on ISSU operation including applicable upgrade paths.

3.6.2.2 Accounting Statistics Alignment

Release 19.10.R1 introduces the **align** configuration flag in **log accounting-policy** entries. If the flag is enabled, the system aligns statistics collection to the next multiple of **collection-interval** in an hour. Such alignment can be used when there is a need to synchronize the statistics collection across different nodes in the network (for example, collect every 20 minutes, aligned to *xx*:00, *xx*:20, *xx*:40).

3.6.2.3 New Counter for Oversized Packets

Release 19.10.R1 introduces a new "Frame too big for port" statistic in the **show card** *slot-number* **fp** *fp-number* **fwd-engine drop-reason statistics** output to track frames that are dropped before transmit when they exceed the MTU of the port.

3.6.2.4 gRPC Certificate Management

Release 19.10.R1 introduces certificate management as a gRPC service. Using the gNOI.CERT RPCs, an external controller can manage TLS certificates on the SR OS node, such as installing new certificates or rotating existing certificates.

3.6.2.5 LAG Port Hash Weight

Release 19.10.R1 introduces the **lag** *lag-id* **hash-weight-threshold** command to control the operational status of the LAG or the IGP link cost based on the sum of the **hash-weight** values for the active links in the LAG.

3.6.2.6 LAG Port Threshold and PXC

Release 19.10.R1 introduces support for LAG **port-threshold** in combination with PXC sub-ports in the same LAG.

3.6.2.7 Hybrid OpenFlow Switch Support for PBF/PBR Redundancy

Release 19.10.R1 introduces the OpenFlow support for redundancy for Layer-2 policy-based forwarding (PBF) and Layer-3 policy-based routing (PBR) with sticky destination. The new OpenFlow experimenter "alu_axn_PBFPBR_Redundancy" allows the programming of a secondary action when the primary destination is down and of the sticky destination hold time.

3.6.2.8 Hybrid OpenFlow Switch Support for IPv6 Addressing for OpenFlow Controller

Release 19.10.R1 adds the support for IPv6 addressing for the OpenFlow controller. The support for an IPv6 address is added to the **config**>**open-flow>of-switch>of-controller** CLI command.

3.6.2.9 Support for SyncE/1588 CPM and CCM Ports

Release 19.10.R1 adds support for the SyncE/1588 ports on the CPM or CCM cards. These ports support receiving a synchronous Ethernet (SyncE) frequency into the central clock and exchanging IEEE 1588 (PTP) messages with external PTP clocks. The ports operate in 100BASE-TX mode only. For PTP, only PTP over Ethernet encapsulation is supported. Support is enabled on the following chassis:

- 7750 SR-7/12/12e
- 7750 SR-7s/14s
- 7950 XRS-20

3.6.2.10 DHCP Client for OOB Management Address and Dual Stack IPv6

In Release 19.10.R1, **autoconfigure** can be configured under **bof** to obtain an IPv6 management address from a DHCP server. This feature is specifically for the out-of-band management port. It should be noted that when **autoconfigure** is enabled under **bof**, a static IP address, static route, or DNS cannot be configured for the respective address family (IPv4 or IPv6). All of these items need to arrive from the DHCP server. **autoconfigure** for IPv4 and IPv6 can be configured to obtain both v4 and v6 addresses and configure a dual-stack management port.

autoconfigure is currently supported on 7750 SR-1s/2s/7s/14s, 7750 SR-1, and 7750 SR-1e.

3.6.2.11 HTTPS URI Support in File Commands

Release 19.10.R1 introduces the support for HTTPS Uniform Resource Identifiers (URIs) for remote files in **file copy**, **type**, **move**, and **delete** commands. If an HTTPS URI redirects to a different URI, the system prompts the operator to accept or refuse redirection. Commands with the **no-redirect** keyword automatically refuse redirection. As with deleting or overwriting files, commands with the **force** keyword will automatically accept redirection.

3.6.2.12 HTTP Proxy Support in File Commands

Release 19.10.R1 introduces the support for HTTP proxy servers in file commands that contain HTTP or HTTPS URIs.

3.6.2.13 IP Filter Match First-Fragment/Not-First-Fragment

Release 19.10.R1 introduces the support for matching the first fragment of a packet, or a fragment that is not the first fragment within ingress IPv4 filters on an FP4-based hardware assembly.

3.6.2.14 LI Configuration Mode Check

In Release 19.10.R1, the **tools**>**perform**>**system**>**management-interface**>**configuration-mode** [**classic**|**mixed**|**model-driven**] **check** [**li**] command is introduced to test if the Lawful Intercept (LI) configuration is ready for a mode switch from classic to either mixed or model-driven mode. The command validates the LI configuration and lists detailed instructions to prepare the LI configuration for a mode switch. To see the list of instructions, operators must have LI access.

3.6.2.15 OSI MAC Layer

Release 19.10.R1 introduces the support for MACsec encryption at Layer 2 (OSI MAC layer). In previous releases, MACsec could only be configured for an entire port. In addition, MACsec could only be signaled using an untagged EAPoLAN packet that was generated and forwarded from that port. There was no support for MACsec to be signaled by a specific VLAN (**encaptype dot1q** or **qinq**). In Release 19.10.R1, MKA (EAPoLAN) packets can be generated using a specific dot1q or QinQ tag. This enables MACsec to be signaled over an Layer-2 VLAN switch network, maintaining and switching the VLAN tags as necessary in the network to connect two or more MACsec-capable routers over the Layer-2 VLAN switch network. In addition, separate VLANs can be used to signal MACsec and use a separate PSK from other customers.

3.6.2.16 ZTP for IPv4 and IPv6 Networks

In Release 19.10.R1, Zero Touch Provisioning (ZTP) can be used to automatically discover a node on the network and download the required files for its provisioning. These files include the latest image, configuration file, and licenses. ZTP uses DHCP to obtain the IPv4/IPv6 address and location of a provisioning file. This provisioning file can be designed by the provider to include information such as the location of the image, configuration file, and licenses. The provisioning file can also include the DNS server to be used to resolve these URLs. The ZTP process will go through these files one by one and download them to the compact flash. ZTP then will reboot the node and the node will use the downloaded files to boot up. ZTP can be used on the out-of-band management port or on the first two slots and first two breakout ports on the first two connectors. ZTP is currently supported on 7750 SR-1s/2s/7s/14s, 7750 SR-1, and 7750 SR-1e.

The correct part number for ZTP should be ordered when ordering the Release 19.10.R1 software license. The part number with ZTP is shipped with a compact flash which includes a bof.cfg file that has a ZTP flag which forces the node to boot up in ZTP mode.

3.6.2.17 LI Support in the MD-CLI and NETCONF

Release 19.10.R1 introduces the support for Lawful Intercept (LI) in the MD-CLI and NETCONF. To switch from classic mode to mixed or model-driven mode when LI is being used, migration steps must be performed. Refer to the *OAM and Diagnostics Guide* to understand the migration steps and key procedures. Also see Known Limitations for more information regarding LI MD-CLI support.

3.6.3 Routing

3.6.3.1 BGP Selective Label-IPv4 Route Installation

Release 19.10.R1 introduces a BGP **selective-label-ipv4-install** configuration option to suppress the programming of received label-unicast IPv4 routes that are not needed to resolve any service endpoint or any BGP next-hop of a BGP VPN route. This option can save IP FIB and label FIB resources in a router that only needs label-IPv4 routes for these purposes.

3.6.3.2 BGP Unresolved Route Leaking from Base to VPRN

Release 19.10.R1 introduces the, **config**>**router**>**bgp**>**next-hop-res**>**allow-unresolved-leaking** command. This command allows BGP routes learned from peers of the Base BGP instance to be leaked to other (VPRN) BGP instances, even if the routes to be leaked do not have a BGP next-hop that can be resolved by the Base instance.

3.6.3.3 BGP Segment Routing Using the Prefix SID Attribute

Release 19.10.R1 introduces the support for BGP Segment Routing (SR). When it is enabled, BGP can send and receive label-unicast IPv4 routes that have an attached prefix SID attribute. The prefix SID attribute advertises an MPLS label index that encodes the network-wide instruction to forward traffic along the ECMP-aware BGP-computed best path(s) to reach the IPv4 prefix in the NLRI field of the route. The BGP prefix SID attribute allows SR to be extended into network domains that may not support an IGP protocol or may not support the SR extensions to an IGP protocol. The BGP prefix SID attribute can also help to create SR paths that transit across multiple administrative domains that do not share IGP SR topology information.

3.6.3.4 BGP Remove-Private Replace Option

Release 19.10.R1 introduces a new **replace** option for the BGP **remove-private** command. When this option is specified, private autonomous-system (AS) numbers in the AS path of advertised BGP routes are replaced with the AS number of the advertising BGP instance, rather than being stripped.

3.6.3.5 Flow Sampling for Tunneled Traffic

Release 19.10.R1 introduces two Cflowd options that allow the inner IP flow information in L2TPv2 or IPv4-over-IPv6-encapsulated traffic to be analyzed. If any of these types of traffic are sampled by Cflowd, and the corresponding option is enabled, Cflowd analysis is performed on the inner IP packet instead of the outer transport headers.

The IPv4-over-IPv6-encapsulated traffic types supported are DS-Lite and MAP-E traffic.

These options are enabled independently using the **config**>**cflowd**>**analyze-l2tp-traffic** and **config**>**cflowd**>**analyze-v4overv6-traffic** CLI commands. The default behavior is for these two options to be disabled.

3.6.3.6 Micro-Loop Avoidance Using Loop-free SR Tunnels (IS-IS)

Release 19.10.R1 enables the micro-loop avoidance feature in the IS-IS instance using Segment-Routing (SR) loop-free tunnels.

The feature is triggered by the receipt of a single P2P link event:

- link addition/restoration
- link removal/failure
- link metric change

IS-IS then runs the main SPF and LFA SPFs and will compute and activate for each SR-ISIS node SID a loop-free SR tunnel applicable to the specific link event. This tunnel acts the micro-loop avoidance primary path for the prefix of the node SID and uses the same outgoing interface as the new computed primary next-hop.

At the expiry of the **fib-delay** timer, IS-IS programs the new primary next-hop(s) for the SR-ISIS node SID tunnel.

This feature is supported with the following contexts:

- IS-IS MT=0 for a SR-ISIS IPv4/IPv6 tunnel (node SID)
- IPv4 and IPv6 SR-TE LSP that use a node SID in their segment list
- IPv4 and IPv6 SR policy that use a node SID in their segment list

3.6.3.7 RFC 6549 OSPFv2 Multi-Instance Extensions

Release 19.10.R1 adds the support for multi-instance extensions (as described in RFC 6549) for OSPFv2 in the Base router. For migration scenarios, OSPFv2 Multi-instance support can be enabled or disabled separately for each instance.

3.6.3.8 Route Policy Action to Suppress BGP Route Installation

Release 19.10.R1 introduces the **disable-route-table-install** route policy action, applicable in BGP peer import policies, that causes matched BGP routes to be suppressed from installation into the route table and tunnel table associated with the BGP instance. The matched BGP routes are still advertisable to other BGP peers, even if next-hop-self applies. Operators should use caution to avoid inadvertent blackholing of traffic.

3.6.3.9 Sticky ECMP for BGP for Addition of BGP Next-hops

Release 19.10.R1 extends the support for sticky ECMP for BGP routes. Prior to Release 19.10.R1, when a BGP route was selected (through policy) for sticky-ECMP, and it was an ECMP route with multiple equal-cost BGP next-hops, then the failure or removal of one or more of these next-hops caused only the affected traffic flows to be re-distributed to the remaining next-hops as evenly as possible. The addition of a new next-hop caused all flows to be redistributed over the new resulting set of next-hops (similar to ECMP without sticky-ECMP configured) causing all flows to be affected, even those using existing next-hops.

Beginning in Release 19.10.R1, when a new next-hop is added for a selected BGP route, the redistribution of traffic is optimised, significantly reducing the number of flows that may be impacted.

3.6.3.10 VPRN BGP Recursive Resolution

Release 19.10.R1 extends the support of BGP-resolving-BGP that is already supported by the Base router BGP instance to VPRN BGP instances. When enabled, a VPRN BGP route can be resolved by another VPRN BGP route (up to four levels of recursion) or by a BGP-VPN route.

3.6.3.11 Weighted ECMP for IGP and BGP Shortcuts over SR-TE Tunnels

Release 19.10.R1 introduces weighted ECMP for IGP shortcuts resolving to next hops over SR-TE tunnels in RTM, BGP labeled unicast routes (RFC 3107) over provisioned SR-TE and RSVP-TE tunnels, BGP unlabeled routes resolving to a static route shortcut over provisioned SR-TE or RSVP LSPs, and 6PE over provisioned SR-TE tunnels.

Weighted ECMP for IGP shortcut tunnels in RTM is enabled using the existing **config**>**router**>**weighted-ecmp** command. Weighted ECMP for BGP shortcuts and labeled routes into SR-TE and RSVP-TE tunnels in TTM is enabled using the **config**>**router**>**bgp**>**next-hop-res**>**weighted-ecmp** command.

3.6.3.12 Policy-based Routing to MPLS Forwarding and SR-TE Policies

Release 19.10.R1 introduces the capability to redirect traffic, using a policy-based routing rule, to either an MPLS forwarding policy or an SR-TE policy. The PBR rule references the policy by its endpoint (for MPLS forwarding policies), or by both its endpoint and color (for SR-TE policies). If the programmed instance of the policy changes, the PBR rule follows this change and redirects traffic to the newly-programmed instance. This PBR rule is applicable on ingress only on Layer-3 interfaces.

3.6.3.13 Static and BGP SR Policies with IPv6 Endpoint

Release 19.10.R1 extends the static and BGP SR policy features with the support of IPv6 endpoint. Release 19.10.R1 introduces the support of the IPv6 AFI for the SR policy NLRI SAFI in the BGP control plane and allows the resolution of the IPv6 next-hop of the following BGP unicast route families:

- Use IPv6 tunnel of type **sr-policy** to resolve the IPv6 next-hop of the following route families: **ipv4**, **ipv6, vpn-ipv4, vpn-ipv6, label-ipv4, label-ipv6, evpn**.
- Use IPv6 tunnel of type **sr-policy** to resolve the IPv4 next-hop of the following route families: **ipv4** and **label-ipv4** (SR policy with **endpoint** 0::0 only).
- Use IPv6 tunnel of type **sr-policy** to resolve the IPv4-mapped IPv6 next-hop of the following route families: **label-ipv6** (SR policy with **endpoint** 0::0 only).

3.6.3.14 PIM Signaling via BIER

Release 19.10.R1 introduces PIM signaling through BIER as a mechanism to signal PIM joins and prunes through a BIER domain with minimum disruption to the PIM domain routers. The ingress BIER edge routers terminates PIM and forwards the PIM joins and prunes through a BIER domain to the egress BIER edge routers. The egress BIER edge routers will track all ingress BIER edge routers that are interested in a particular (S,G). From a datapath point of view, the multicast PDUs are encapsulated in BIER headers and forwarded to all ingress BIER edge routers that are interested in that (S,G). The BIER header is removed, and the multicast PDU is forwarded to the PIM domain.

3.6.3.15 Route Policy Support of Named Entries

Release 19.10.R1 introduces the support for named policy entries. Prior to Release 19.10.R1, only numbered policy entries have been supported, where each entry is given a sequence number and the rules are always evaluated in order of ascending sequence number. With named entries, each entry is a specific string-format name up to 255 characters in length, and entries are evaluated in user-order (that is, the order they appear in the configuration). This feature is accompanied by new MD-CLI and NETCONF infrastructure that allows the order of user-ordered list entries to be rearranged using **insert** commands. A single policy-statement cannot have a mix of named and numbered entries. It is not possible to have any policy-statements with named entries in classic or mixed mode. This new feature is only supported in **model-driven management-interface configuration-mode**.

3.6.3.16 Support for Strict IS-IS Weighted ECMP

Release 19.10.R1 introduces the **weighted-ecmp** command that enables weighted load-balancing for IS-IS. This command assumes that each member of the weighted ECMP bundle has an IS-IS interface weight configured. If one or more **weighted-ecmp** members does not have an interface weight configured, then the bundle will fall back to ECMP from **weighted-ecmp**. Beginning in Release 19.10.R1, the **weighted-ecmp** command has been extended with a new configuration **weighted-ecmp strict** option, forcing each member of the **weighted-ecmp** bundle to have an explicit configured IS-IS interface weight before the member interface is taken into operation within the **weighted-ecmp** bundle. This avoids fallback to ECMP when no explicit IS-IS weight was configured.

3.6.3.17 IPv6 Traffic Engineering (IS-IS)

Release 19.10.R1 extends the traffic engineering capability in IS-IS with the support of IPv6 TE links and nodes.

IS-IS, BGP-LS and the TE database are enhanced with the additional IPv6 link TLVs and TE link TLVs and provides three modes of operation of the IPv4 and IPv6 traffic engineering in a network.

- Legacy Mode: enables the existing traffic engineering behavior for IPv4 RSVP-TE and IPv4 SR-TE. Only the RSVP-TE attributes are advertised in the legacy TE TLVs to be used by both RSVP-TE and SR-TE LSP path computation in the TE domain routers. In addition, IPv6 SR-TE LSP path computation can now use these common attributes.
- Legacy Mode with Application Indication: This mode is intended for cases where link TE attributes are common to RSVP-TE and SR-TE applications and have the same value, but the user wants to indicate on a per-link basis which application is enabled. Routers in the TE domain will use these attributes to compute a path for IP4 RSVP-TE LSP and IPv4/IPv6 SR-TE LSP.
- Application-Specific Mode: The application-specific mode of operation is intended for future use cases where TE attributes have different values in RSVP-TE and SR-TE applications or are specific to one application (e.g., Unreserved bandwidth attribute).

Routers in the TE domain will use these attributes to compute the path for IP4 RSVP-TE LSP and IPv4/IPv6 SR-TE LSP.

3.6.3.18 LFA Policy Support with base LFA, remote LFA, and TI-LFA backup paths (SR-ISIS)

Release 19.10.R1 extends the LFA policy feature to all Loop-Free Alternate (LFA) methods in Segment Routing (SR).

The LFA policy feature provides the user with policy control of the LFA backup next-hop selection within Shortest Path First (SPF) calculation in IGP. The feature introduces the concept of route next-hop policy template to influence LFA backup next-hop selection. The template supports the following policy attributes:

- IP Admin-Group include/exclude constraint
- IP Shared Risk Loss Group (SRLG) constraint
- Protection type preference: link or node protection
- Next-hop type preference: IP or tunneled

The policy is applied to the interface of the primary next-hop of the destination prefix of the following SR tunnel types:

- IPv4 and IPv6 SR-ISIS node SID tunnels

This feature is not supported with IPv4 and IPv6 SR-ISIS adjacency SID tunnels.

These LFA policy attributes are checked against the outgoing interface used by the LFA/RLFA/TI-LFA backup path and selects the backup paths which satisfy the policy attributes.

The LFA policy indirectly applies to an IPv4 or IPv6 SR-TE LSP, and to an IPv4 or IPv6 SR policy, which use any of the above SR tunnels as the top SID in their SID list.

The LFA policy also indirectly applies to IPv4 LDP FECs when the LDP **fast-reroute backup-sr-tunnel** option is enabled and the FEC is protected with a SR tunnel.

3.6.4 MPLS

3.6.4.1 Configure SID for Path Hops

Release 19.10.R1 introduces the ability to configure an explicit SID value for hops of an SR-TE LSP path. This explicit SID value is an MPLS label for a SID along the path of the LSP. The SID value is configured using the new **config**>**router**>**mpls**>**path**>**hop** *index* **sid-label** command.

3.6.4.2 LDP Remote LFA for IPv4 FECs

Release 19.10.R1 introduces the capability for LDP to setup Remote LFA backups to protect primary traffic using the **config**>**router**>**isis**>**lfa**>**augment-route-table**. LDP Remote LFA applies to IPv4 FECs. LDP Remote LFA requires the targeted sessions (between source node and PQ node) to be manually configured and is available with IS-IS only. LDP remote LFA is designed to be operated in LDP only environments; as such, LDP does not establish a Remote LFA backup when in the presence of LDP-over-RSVP-TE or LDP over SR-TE tunnels.

3.6.4.3 LSP Self Ping

Release 19.10.R1 introduces LSP self ping for RSVP-TE LSPs. LSP self ping is a lightweight mechanism to check that an LSP data path has been programmed for all LSRs following the receipt of an RESV for the path before switching traffic to it. When enabled, LSP self ping packets are periodically sent until a response is received or configured duration, on LSP candidate paths when an LSP is about to switch the active path or is undergoing an MBB. It is only used in cases where the previously active path remains up.

LSP self ping is enabled for all RSVP-TE LSPs using the **config**>**router**>**mpls**>**lsp-self-ping**>**rsvp-te enable** command. It may be explicitly disabled for a specific RSVP LSP or auto-LSP using the **lsp-self-ping disable** command under the LSP context or the LSP template.

LSP self ping uses UDP port 8203.

3.6.4.4 LDP ECMP Next-hop Limit Increased to 64 (Unicast)

Release 19.10.R1 increases the maximum number of ECMP next hops available to LDP from 32 to 64. This is applicable to cases where LDP resolves a FEC over network IP interfaces, IGP shortcut tunnels with RSVP-TE and SR-TE LSPs, LDP-over-RSVP, and for FEC 129 and P2MP FECs.

The **config**>**router**>**ldp**>**max-ecmp-routes** command is introduced to configure the increased limit, subject to a maximum set in the **config**>**router**>**ecmp** context. The default limit for LDP remains at 32.

In Release 19.10.R1, this increase is introduced for Unicast only.

3.6.4.5 SRLG for SR-TE Secondary Paths

Release 19.10.R1 enables the **srlg** command for secondary paths for SR-TE LSPs. When configured with CSPF, the router attempts to calculate a path for the secondary that is SRLG divert from the primary.

3.6.4.6 SR-Policy Traffic Statistics

Release 19.10.R1 introduces the ability to configure the collection of traffic statistics on the egress data-path of ingress LERs of SR policies. The collected statistics can be accessed using the CLI (using **show** or **monitor** outputs) and via telemetry. Statistics collection is enabled globally for SR policies and applies to both static and signaled SR-policies. Traffic statistics are collected per-segment list. However, the system will allocate at most 32 statistic indexes for all the instances of a specific policy. Traffic statistics are provided without forwarding class or QoS profile distinction.

3.6.4.7 SR-TE LSP Traffic Statistics

Release 19.10.R1 introduces the ability to configure the collection of traffic statistics on the egress data-path of ingress LERs of SR-TE LSPs. The collected statistics can be accessed by using the CLI (**show** or **monitor**) and via Telemetry. Statistics are collected for each path (primary, and backup(s) if any) of the SR-TE LSP. Traffic statistics are retained across switchover except for secondary non-standby paths. Traffic statistics are provided without forwarding class or QoS profile distinction.

3.6.4.8 IGP SIDs Traffic Statistics

Release 19.10.R1 introduces the capability to collect traffic statistics for IGP SIDs (Node SIDs, Adjacency SIDs, and Adjacency Set SIDs) on both ingress and egress. Traffic statistics are collected regardless of the forwarding class or the QoS profile. Release 19.10.R1 also provides a means to display, monitor, or clear the traffic counters, as well as access these via telemetry.

3.6.4.9 IPv6 SR-TE LSP

Release 19.10.R1 adds the support of IPv6 destinations to the SR-TE LSP configuration. It also extends the MPLS path configuration with hop indices that include IPv6 addresses.

IPv6 SR-TE LSP is supported with the hop-to-label path computation method only and requires the enabling of the IPv6 traffic engineering feature in IS-IS.

The IPv6 SR-TE LSP is supported in the following routing and service contexts:

- VPRN auto-bind (families vpn-IPv4, vpn-IPv6)
- BGP shortcut and BGP-LU (families IPv4, IPv6, label-IPv4, label-IPv6)
- BGP EVPN (family EVPN)
- Static Route with IPv6 next-hop

3.6.4.10 Support of Forwarding CPM-originated Packets over LFA Backup

Release 19.10.R1 introduces the ability for CPM-originated packets to be forwarded over an LFA backup path when activated on the router. The LFA backup can be of type IP or tunnel and can be computed by base LFA, remote LFA, or Topology-Independent LFA (TI-LFA).

CPM-originated packets are IP and MPLS OAM packets, BFD and seamless BFD (sBFD) packets, and routing and MPLS control-plane protocol packets when forwarded to a destination reachable via an IP route or a tunnel which activated a LFA backup at the node originating the packet. The destination may be reachable via:

- an IP next-hop
- an IGP shortcut (RSVP-TE/SR-TE)
- an LDP, SR-ISIS, or SR-OSPF tunnel
- an SR-TE LSP
- an SR policy
- BGP-LU tunnel when resolved over one or a hierarchy of any of the above

This feature does not support the following CPM packets:

- IP BFD packets
- LSP-trace packets
- Nokia proprietary OAM packets which do not perform an IP route lookup. An example would be **sdp-ping** and **sdp-keepalive**.

3.6.5 Services

3.6.5.1 Affect LAG Signaling on EVPN Single-active Multi-homing State

Release 19.10.R1 introduces the ability to use LACP or Power-off LAG standby signaling in an EVPN single-active multi-homing scenario so the CE does not send traffic to the non-Designated Forwarder (non-DF) PE. To use LAG standby signaling on the Ethernet Segment (ES), an **oper-group** can be now added under the **ethernet-segment** context. If all the SAPs contained in the ES are operationally down due to the non-DF state, the ES **oper-group** will go down. The LAG associated with the ES then signals “standby” state to the CE, if the **monitor-oper-group** (for the ES **oper-group**) is configured on the LAG. Upon failure on the DF PE, as soon as the first ES SAP comes up in the non-DF PE, the ES **oper-group** comes up and the local LAG no longer signals standby state.

This feature requires that the same PE is DF for all of the ES SAPs, and there is a single ES in the LAG. Therefore, the feature is not supported with DF Election algorithm **auto**, **evi**, or **isid** ranges, or virtual ESs. In addition, ES **oper-groups** cannot be added to any other objects.

3.6.5.2 EVPN-VXLAN Network-interconnect Multi-homing (in Two VXLAN Instance Services)

Release 19.10.R1 completes the EVPN-VXLAN multi-homing support by adding the support for network-interconnect VXLAN Ethernet Segments (ESs) in services with two VXLAN instances. Prior to Release 19.10.R1, the redundancy for dual VXLAN instance services was based on anycast redundancy. In particular, the following scenarios are now supported:

- Network-interconnect VXLAN ES in dual instance VPLS/R-VPLS services with two EVPN-VXLAN instances
- Network-interconnect VXLAN ES in dual instance VPLS/R-VPLS with one EVPN-VXLAN instance and one static VXLAN instance

With the support for network-interconnect VXLAN ESs, local SAPs are now supported on dual VXLAN instance services.

3.6.5.3 ECMP in EVPN Services

Release 19.10.R1 introduces a number of ECMP features in EVPN Services.

ECMP for Layer-2 Unicast traffic on Epipe and VPLS services for EVPN-MPLS destinations:

- This is enabled using the **config**>**service**>**epipe|vpls**>**bgp-evpn**>**mpls**>**auto-bind-tunnel# ecmp** *number* commands and allows the resolution of an EVPN-MPLS next-hop to a group of ECMP tunnels of type RSVP-TE, SR-TE or BGP.

ECMP for Layer-3 Unicast traffic on R-VPLS services with EVPN-MPLS destinations:

- This is enabled using the **vpls**>**bgp-evpn**>**mpls**>**auto-bind-tunnel**>**ecmp** and **vpls**>**allow-ip-int-bind**>**evpn-mpls-ecmp** commands.
- The VPRN Unicast traffic (IPv4 and IPv6) will be sprayed among *m* paths, *m* being the lowest value of (16,*n*), where *n* is the number of ECMP paths configured in the **auto-bind-tunnel**>**ecmp** command.
- CPM-originated traffic will not be sprayed and will pick up the first tunnel in the set.
- This feature is limited to FP3 and higher systems.

ECMP for Layer-3 multicast traffic on R-VPLS services with EVPN-MPLS or EVPN-VXLAN destinations:

- For EVPN-MPLS destinations, this is enabled using the **vpls**>**allow-ip-int-bind**>**ip-multicast-ecmp** and **auto-bind-tunnel**>**ecmp** commands. The VPRN multicast traffic (IPv4 and IPv6) will be sprayed among up to *m* paths, *m* being the lowest value of (16,*n*), where *n* is the number of ECMP paths configured in the **auto-bind-tunnel**>**ecmp** command.
- For EVPN-VXLAN destinations, this is enabled using the **vpls**>**allow-ip-int-bind**>**ip-multicast-ecmp** and **router**>**ecmp** commands. The VPRN multicast traffic (IPv4 and IPv6) will be sprayed among up to *m* paths, *m* being the lowest value of (16,*n*), where *n* is the number of ECMP paths configured in the **router**>**ecmp** command.

3.6.5.4 IPv6 Tunnel Resolution for EVPN MPLS Services

Release 19.10.R1 introduces the support for IPv6-based tunnel types for EVPN MPLS auto-bind resolution. The following IPv6 tunnels are supported:

- IPv6 SR-ISIS

- IPv6 SR-OSPFv3
- IPv6 SR-TE
- IPv6 SR policy
- IPv6 RIB API
- IPv6 MPLS forwarding policy
- IPv6 LDP

These tunnel types are supported for auto-bind resolution and behave as any other MPLS tunnel type as far as EVPN is concerned. In particular, they are supported in all EVPN services that support MPLS tunnels (for example, VPLS, Epipe, R-VPLS, B-VPLS, and dual-BGP-instance services).

The **vpls|epipe**>**bgp-evpn**>**mpls**>**route-next-hop** *ip-address* command must be configured with the IPv6 address that the system will use as a next-hop for the EVPN service routes.

See Limited Support Features for more information.

3.6.5.5 EVPN Multi-Homing Support for MPLS Tunnels using IPv6 or Non-system IPv4 Addresses

Release 19.10.R1 introduces the support for EVPN multi-homing on PEs that use non-system IPv4 or IPv6 addresses as BGP next-hops.

The **vpls|epipe**>**bgp-evpn**>**mpls**>**route-next-hop** *ip-address* command must be configured with the non-system IPv4 or IPv6 address that will be used as next-hop for the EVPN routes for the service (MAC/IP routes, Inclusive Multicast Ethernet Tag routes, and auto-discovery per EVI routes). In addition, when multi-homing is used in the service, the same *ip-address* that is configured at service level must be configured in the Ethernet Segment **es**>**route-next-hop** *ip-address* and **es**>**es-orig-ip** *ip-address* commands. By configuring the same *ip-address*, the designated forwarder (DF) candidate list will be built with the correct IP addresses and DF election can be performed.

The **vpls|epipe**>**bgp-evpn**>**mpls**>**route-next-hop** cannot be configured with a value different than *system-ipv4* address in EVPN services using IGMP-snooping, PIM-snooping, or provider-tunnel.

See Limited Support Features for more information.

3.6.5.6 Operational Groups for EVPN Services

Release 19.10.R1 introduces the **oper-group** command that is supported under the **bgp-evpn**>**mpls** context for Epipe, VPLS, R-VPLS, and B-VPLS services, and under **bgp-evpn**>**vxlan** for VPLS and R-VPLS services (not Epipe). When using an **oper-group** under **bgp-evpn**:

- The **oper-group** cannot be configured under any other object within the same or different service.
- The operational status "up" of an **oper-group** configured on a BGP-EVPN instance is determined by the existence of at least one EVPN destination in the instance.

The operational group goes down in the event of the following:

- Service admin-state disable
- BGP-EVPN VXLAN or MPLS admin-state disable
- Removal of all of the EVPN destinations in the EVPN instance

3.6.5.7 Weighted ECMP

Release 19.10.R1 introduces the support for weighted ECMP for the following objects:

- Epipe and Ipipe VLL spoke-SDP termination on IES and VPRN interfaces over RSVP-TE tunnels
- Epipe spoke-SDP termination on the VPLS component of an R-VPLS over RSVP-TE tunnels

3.6.5.8 VSI-Import and Export Policies on Epipes (VPWS and EVPN)

Release 19.10.R1 adds the support for **vsi-export** and **vsi-import** policies in Epipe services under the **config**>**service**>**epipe**>**bgp** context. These policies can be used by BGP EVPN, BGP VPWS, and BGP multi-homing in Epipe services. Up to five policies of each type (export and import) can be configured. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy to match is applied.

3.6.5.9 Control Word Support on PW-Port

Release 19.10.R1 adds the support for control words (CW) on:

- a spoke-SDP or BGP EVPN associated with FPE-based PW-ports
- a spoke-SDP associated with a fixed PW port

3.6.5.10 Control Word Support on Mirror Destination Spoke-SDPs

Release 19.10.R1 adds the support for the PW control word to mirror destination spoke-SDPs and are a part of a mirror service of type **ether**. The **control-word** command can now be configured under the **config**>**mirror**>**mirror-dest**>**spoke-sdp** and **config**>**mirror**>**mirror-dest**>**remote-source**>**spoke-sdp** contexts.

3.6.5.11 EVPN IPv4 Host Mobility within the Same R-VPLS Service

Release 19.10.R1 introduces the support for EVPN host mobility as stated in section 4 of *draft-ietf-bess-evpn-inter-subnet-forwarding*. This feature is enabled using the following commands that are configured on the R-VPLS interface of the PEs between which the IPv4 host is moving, for example, PE1 and PE2:

- **config**>**service**>**vprn**>**if**>**arp-host-route**>**populate dynamic**
- **config**>**service**>**vprn**>**if**>**vpls**>**evpn**>**arp**>**no learn-dynamic**
- **config**>**service**>**vprn**>**if**>**vpls**>**evpn**>**arp**>**advertise dynamic**

When host-1 is attached to PE1's R-VPLS with the above settings and sends an ARP packet, PE1 creates an ARP entry for host-1, a host route in the route-table (due to the **populate dynamic**) and advertises the ARP entry in an EVPN MAC/IP route in the context of the R-VPLS (due to the **advertise dynamic**). Since PE2 is attached to the same R-VPLS, it may receive the ARP packet on an EVPN tunnel, but PE2 only learns host-1's ARP entry from the EVPN route (due to **no learn-dynamic**). Prior to Release 19.10.R1, PE2 learned the ARP entry via EVPN tunnel and created a host-route, even if host-1 was not directly attached to PE2.

Also, with the above commands and following Section 4 of *draft-ietf-bess-evpn-inter-subnet-forwarding*, when host-1 moves to PE2, PE1 removes its corresponding ARP entry and host route, while PE2 adds a new dynamic ARP entry, creates a host route, and advertises the MAC/IP route for host-1's MAC and IP.

Along with this feature, the **config**>**service**>**vprn**>**if**>**vpls**>**evpn**>**arp**>[**no**] **flood-garp-and-unknown-req** command has been added to control the flooding of CPM-generated ARP messages on EVPN destinations. The **no** form of the command prevents ARP flooding to EVPN destinations.

3.6.5.12 EVPN-MPLS Layer-3 OISM

Release 19.10.R1 introduces the support for EVPN-MPLS Optimized Inter-Subnet Multicast (OISM) based on *draft-ietf-evpn-irb-mcast*. The solution is supported along with EVPN-MPLS destinations on R-VPLS services, ingress replication and IPv4 multicast traffic. OISM provides an efficient Layer-3 multicast forwarding solution in EVPN networks. Aspects of the OISM support in Release 19.10.R1 are:

- Receivers use IGMP or static IGMP groups to pull the multicast traffic from the PEs.
- PEs use the EVPN route type 6 or Selective Multicast Ethernet Tag (SMET) route to pull multicast traffic from ingress PEs.
- PEs attached to the same OISM tenant will have one or more “ordinary” R-VPLS services and one Supplementary Broadcast Domain (SBD) R-VPLS for the tenant VPRN.
  - The SBD is an R-VPLS service configured as **evpn-tunnel supplementary-broadcast-domain** and with **bgp-evpn**>**sel-mcast-advertisement**. It is used to aggregate and advertise the SMET routes to pull the multicast traffic, and also to receive the multicast traffic if the PE is not attached to the source ordinary R-VPLS.
  - The ordinary R-VPLS services provide connectivity to the sources and receivers and use regular IGMP procedures with the sources and receivers.
  - The VPRN attached to ordinary and SBD R-VPLS services must be configured with PIM on all R-VPLS interfaces and IGMP on the ordinary R-VPLS interfaces. The SBD interface must be configured with **pim**>**interface**>**multicast-sender always** so that the received multicast traffic can pass the RPF checks.
- The receivers and sources can use EVPN multi-homing to connect to the OISM PEs. When there are only two PEs attached to the Ethernet Segment (ES), the receiver IGMP-snooping multicast state can be synchronized via Multi-Chassis Synchronization protocol (MCS).

3.6.5.13 GRE Fragmentation and Reassembly

Release 19.10.R1 provides fragmentation and reassembly support for Layer-2 services using GRE transport. Specifically, those services with a pseudo-wire template configured with **auto-gre-sdp** now have the option for packets to be fragmented when transmitted out of the network interface if the MTU of the packet on SAP ingress is larger than the service MTU. The reassembly function is provided by the **isa-bb** or **isa2-bb** application where GRE transport fragments are received on the terminating node and reassembled before further network ingress processing is performed on the packet.

This feature can be enabled to work with the GRE termination on interface IP address feature.

3.6.5.14 GRE Termination on Interface IP Address

Release 19.10.R1 provides the ability to terminate services using GRE transport on the network interface IP address. Terminating services with GRE transport on a network interface address is an alternative option to using the system IP address. This may be useful for deployments where the node is connected to two service providers (possibly in a redundancy scheme) using a network interface to each service provider, and the system IP address cannot be used for routing, only the interface IP address from each provider.

3.6.6 Subscriber Management

3.6.6.1 Data-Triggered ESM Host Mobility

Release 19.10.R1 supports ESM subscriber host mobility to allow the replacement of existing hosts by mobility SHCV events in following cases that occur between the SAPs within the same group interface or in the different group interfaces:

- Protocol-triggered ESM host on the new SAP replaces protocol-triggered ESM host on the old SAP
- Protocol-triggered ESM host on the new SAP replaces data-triggered ESM host on the old SAP
- Data-triggered ESM host on the new SAP replaces data-triggered ESM host on the old SAP

- Data-triggered ESM host on the new SAP replaces protocol-triggered ESM host on the old SAP

3.6.6.2 Multi-Chassis Synchronization Support for Usage-Monitoring (Gx)

Release 19.10.R1 introduces the synchronization of volume-based Gx usage-monitoring in an SR OS dual-homed environment. Usage counters are periodically synchronized between the chassis, with a configurable synchronization interval. This functionality relies on ESM and Diameter dual-homing, and is supported only on the new Diameter base.

3.6.6.3 Diameter Multi-Chassis Redundancy on New Base

Release 19.10.R1 introduces Diameter multi-chassis redundancy for the Diameter base that was introduced in Release 16.0. The older Diameter base implementation relies on proxy-based dual-chassis redundancy. Operators are urged to transition to the new Diameter base implementation. The older Diameter base implementation and the proxy-based multi-chassis redundancy model are in maintenance mode only, without any further feature enhancement planned.

3.6.6.4 Multi-Chassis Synchronization of RADIUS Usage Counters

Release 19.10.R1 introduces the support for synchronizing usage counters that can be reported through RADIUS accounting in a dual-homed BNG scenario. The master SRRP node keeps the total amount of the statistics that are to be reported. The master synchronizes those statistics in regular intervals via MCS to the standby node. In this fashion, the master copy of the total statistics is maintained on both nodes and failure cases of link, node, and so on, can be recovered from the surviving node.

3.6.7 Application Assurance

3.6.7.1 Flow Attributes

Release 19.10.R1 introduces flow attributes that provide per-flow additional metadata extracted from traffic by AA stateful flow processing and which complement AA Application and **app-group** classification. Flow attributes can be used for analytics and control use cases. The flow attributes supported in Release 19.10.R1 are abr_service, audio, download, encrypted, real_time_communication, upload, and video.

3.6.7.2 Web-service URL Classification

In Release 19.10.R1, AA introduces a web-service for URL classification into content and threat categories that may be used for parental control services, content based filter lists such as the IWF filter list, and to block malicious websites. AA sends an API call to the configured web-service containing the URL that a user wants to access, and the web-service provides the content category of the URL (for example, “Shopping”).

Depending on the web-service profile activated for the user, AA decides on the action to perform; that is, allow the user to access the page or block access and redirect the user to an informative page.

The operator can configure up to eight different profiles that define which categories should be blocked. A profile can be dynamically associated to a subscriber by the AAA/PCRF.

With the API-based parental control feature, the operator does not have to configure filter policies in any external element since AA makes all policy decisions based on the classification (URL category) received from the web-service.

3.6.7.3 NLB-DEM

In Release 19.10.R1, Non-Location-Based Dynamic Experience Management (NLB-DEM) can detect access congestion events and trigger a configured subscriber policing override for AA subscribers. When a subscriber is declared to be in a congestion state, the per-subscriber congestion policer rates are triggered. This overrides any pre-existing per-subscriber policer rates, including time-of-day policer rates. These per-subscriber congestion policer rates are applied for the duration of time that the subscriber is in a congestion state.

3.6.8 OAM

3.6.8.1 MPLS OAM Support in IPv4/IPv6 SR Policy

Release 19.10.R1 extends the support for **lsp-ping**, **lsp-trace**, and ICMP tunneling probes to IPv4 and IPv6 SR policies.

- **lsp-ping** provides the ability to test the path to the endpoint of the policy. Each SID list represents a different SR-TE path and as such the feature option allows the user to select to which SID list to direct the MPLS echo request message. Furthermore, the path-destination option allows the echo request message to be directed to a specific next-hop within the selected segment list.
- **lsp-trace** provides the ability to discover the various ECMP paths of each SID list in the policy via the use of the DDMAP TLV and the path-destination option.
- ICMP tunneling allows any router in the path of the SR policy to return the incoming label stack in the ICMP reply message to a UDP traceroute packet.

3.6.8.2 MPLS OAM Support in IPv6 SR-OSPF3 Tunnel

Release 19.10.R1 extends the support of lsp-ping, lsp-trace, and ICMP tunneling probes to IPv6 SR-OSPF3 tunnels.

3.6.9 VSR

3.6.9.1 Unnumbered Subscriber Interfaces

Release 19.10.R1 introduces the support for IPv4 unnumbered subscriber interfaces (**unnumbered** and **allow-unmatching-subnets**) and IPv6 unnumbered subscriber interfaces (**allow-unmatching-prefixes**) on VSR.

3.6.9.2 NGE for L2 auto-gre-sdp Services

Release 19.10.R1 introduces a third type of MPLS services encryption approach for auto-bind Layer-2 services. In prior releases, network group encryption (NGE) of MPLS-based services included enabling NGE on VPRN services and manually configured SDPs carrying a variety of services using those SDPs. In Release 19.10.R1, NGE for Layer-2 **auto-gre-sdp** services adds encryption-keygroup inbound/outbound to PW templates, enabling NGE for BGP-VPLS and BGP-VPWS-based services using **auto-gre-sdp**. The NSP NFM-P is integral in managing NGE-enabled MPLS services. Refer to the NSP NFM-P documentation set and release notes for detailed configuration notes, and availability of managing NGE on Layer-2-**auto-gre-sdp** based services using the NSP.

3.7 Release 19.7.R2

There are no new major features in Release 19.7.R2.

3.8 Release 19.7.R1

3.8.1 Hardware

3.8.1.1 7950 XRS 12-port Universal QSFP-DD XMA

Release 19.7.R1 introduces the 12-port Universal QSFP-DD XMA for the 7950 XRS platforms and offers 2.4T FD with intelligent fan-in/fan-out up to 4T. The XMA supports QSFP28, QSFP+, and QSFP-DD optical modules with the ability to mix-and-match any port, speed, or configuration up to 400G. The XMA is available as three Pay-As-You-Grow hardware capacities:

- 8-port 1.6T
- 12-port 2.4T
- 12-port 2.4T and intelligent fan-in/fan-out up to 4T

Each Pay-As-You-Grow hardware capacity is available in CR (Core Router), ER (Edge Router), and HE (High-Scale Edge Route) functional variants.

Both XMA variants, the 12-port 2.4T QSFP-DD and the 12-port 2.4T QSFP-DD with intelligent fan-in/fan-out up to 4T, operate as an 8-port 1.6T XMA variants when installed in the 7950 XRS-20 and offer full functionality when installed in the 7950 XRS-20e.

Upgrade licenses are available to increase the line rate and unlock connectors from 1.6T up to 2.4T with 4T of intelligent fan-in/fan-out and upgrade from CR to ER and HE.

3.8.1.2 7950 XRS 6-port CFP2-DCO XMA

Release 19.7.R1 introduces the 6-port CFP2-DCO XMA for the 7950 XRS platforms, with each of the six ports supporting CFP2 and CFP2-DCO 1x100G and 2x100G (2x100G client side only) coherent optics in a dual-rate form factor. This XMA offers 1.2T FD line rate in the 7950 XRS-20e and 800G FD line rate in the 7950 XRS-20 with intelligent fan-in/fan-out up to 1.2T. It is available in CR, ER, and HE functional variants with upgrade license from CR to ER and ER to HE. This XMA supports IEEE 1588 Port-Based Timestamping (PBT).

3.8.1.3 APEQ-HVDC-4400

Release 19.7.R1 introduces the High-Voltage DC Advanced Power Equalizer Module (APEQ) APEQ-HVDC-4400 for use on the 7950 XRS-20/20e. At 4400W, this APEQ provides almost 50% more power than the current APEQ-HVDC-3000. Offering both N+1 power redundancy and feed redundancy, the APEQ enables complete powering capabilities to fully-configured FP4 systems.

3.8.1.4 7750 SR 2-port 100G QSFP28 Multi-Service MDA-e

Release 19.7.R1 supports 2-port 100G QSFP28 Multi-Service MDA-e on 7750 SR IOM4-e, IOM4-e-B, and IOM4-e-HS, and 7750 SR-e chassis. Key features of this MDA-e are:

- 100G FD
- MACsec support on all ports
- 4x10GE, 1x40GE (QSFP+, LAN)
- 4x25G Breakout
- 100GE (QSFP28) + RS-FEC (Clause 91 FEC)

3.8.1.5 2x100G Support for 7750 SR 3-port CFP2-DCO MDA-e-XP

Release 19.7.R1 introduces the support for 2x100G when using the CFP2-DCO pluggable optic on the 7750 SR 3-port CFP2-DCO MDA-e-XP. To unlock these two 100G ports, the connector should be configured with a breakout of **c2-100g**.

3.8.2 System

3.8.2.1 Ethernet Satellite Client-Down-Delay Support

Release 19.7.R1 introduces a new **config**>**system**>**satellite**>**eth-sat**>**client-down-delay** *seconds* configuration parameter to allow an Ethernet satellite client port to be operationally brought down if no uplink is available. The new command specifies the delay between when the last configured uplink for a client port goes down, and when the client port should be brought down. The range for this timer is a value of five (5) to 1800 seconds.

3.8.2.2 ZTP for IPv4 Networks

Release 19.7.R1 introduces Zero Touch Provisioning (ZTP) for IPv4 networks supported via out-of-band management port and in-band interfaces on the first two slots and breakout ports on the first two connector ports of the 7750 SR-1 and 7750 SR-1s/2s/7s/14s. ZTP auto-discovers the node IP address via DHCP and download the image and the configuration of the operator file server. The correct part number for Operating Software License with ZTP must to be ordered for the node to boot in ZTP mode.

3.8.2.3 HTTP URI Support in File Commands

Release 19.7.R1 introduces the support for HTTP URIs for remote files in **file copy**, **type**, **move**, and **delete** commands. If an HTTP URI redirects to a different URI, the system prompts the operator to accept or refuse redirection. Commands with the **no-redirect** keyword automatically refuse redirection. As with deleting or overwriting files, commands with the **force** keyword will automatically accept redirection.

3.8.2.4 DHCP Client for OOB IPv4 Management Address

In Release 19.7.R1, **autoconfigure** can be configured under **bof** to obtain an IPv4 management address from a DHCP server. This feature is specifically for the out-of-band management port. It should be noted that when the **autoconfigure** is enabled under **bof**, a static IP address, static route, or DNS cannot be configured. All of these items need to arrive from the DHCP server.

3.8.3 Routing

3.8.3.1 BGP Default Route Origination

Release 19.7.R1 adds the **send-default** command to BGP configuration contexts in the Base router and VPRNs. This new command causes an artificially-generated default route (for IPv4, IPv6, or both) to be advertised to peers within the scope of the command. The generated default is advertised even if the local router does not have any default routes installed in its FIB.

3.8.3.2 BGP AS-Override in Base Context

Release 19.7.R1 introduces the support for BGP AS-override functionality in the Base router BGP context. Prior to Release 19.7.R1, BGP AS-override was only supported in VPRN BGP instances.

3.8.3.3 BGP Convergence - Delayed Route Advertisement

Release 19.7.R1 introduces **min-wait-to-advertise** and **max-wait-to-advertise** BGP configuration options to define how BGP reconverges after the router or the BGP instance restarts. With these changes, BGP can be instructed to delay its route advertisement (RIB-OUT) processing until all peers that reconnected quickly enough after the last restart have sent all their routes. This option is only available for IPv4-unicast and IPv6-unicast routes belonging to the Base router BGP instance or to a VPRN BGP instance.

3.8.3.4 Clear BGP Sessions by Sending Route Refresh Message

Release 19.7.R1 introduces the support for a **soft-route-refresh** option in the **clear router bgp neighbor** command. This option keeps a session and sends one or more ROUTE_REFRESH messages to the peer, each requesting that the peer resend all RIB-OUT routes for a specific address family. This tool can be used to debug and troubleshoot route advertisement issues.

3.8.3.5 Improved Handling of BGP MED in Route Policy Actions

Release 19.7.R1 introduces route policy action commands to manipulate the MED attribute of BGP routes. The MED value can be set independently from the IGP metric value if a route policy entry matches both BGP and IGP routes. When MED is configured to track the route or tunnel cost to reach the BGP next-hop, the new **min-igp** option advertises the lower-bound cost and avoids excessive MED churn. A new expression format allows for more sophisticated adjustments to the advertised MED than in previous releases which only supported simple addition or subtraction operations.

3.8.3.6 IS-IS Minimum Remaining Lifetime

Release 19.7.R1 introduces the **lsp-minimum-remaining-lifetime** command in the **config**>**router**>**isis** and **config**>**service**>**vprn**>**isis** contexts. This command configures the minimum value to which the remaining lifetime of the LSP is set. When using IS-IS as IGP, it is possible for a corruption of the Remaining Lifetime field in an LSP to go undetected and cause premature aging of the LSP. Undetected Remaining Lifetime corruption causing premature aging of an LSP in IS-IS can be abused as a potential denial-of-service attack, or may cause perpetual LSP flooding storms caused by ongoing LSP corruption. This command can be used to avoid flooding storms of a denial-of-service attack.

3.8.3.7 LFA Policy Support with Base LFA, Remote LFA, and TI-LFA Backup Paths (SR-OSPF)

Release 19.7.R1 extends the LFA policy feature to all the Loop-Free Alternate (LFA) methods in Segment Routing (SR).

The LFA policy provides policy control of the LFA backup next-hop selection within Shortest Path First (SPF) calculation in IGP. The feature introduces the concept of route next-hop policy template to influence LFA backup next-hop selection. The template supports the following policy attributes:

- IP Admin-Group include/exclude constraint
- IP Shared Risk Link Group (SRLG) constraint
- Protection type preference: link or node protection
- Next-hop type preference: IP or tunneled

The policy is applied to the interface of the primary next-hop of the destination prefix of the following SR tunnel types:

- IPv4 SR-OSPF node SID tunnels
- IPv4 SR-OSPF backup node SID tunnels

The feature is not supported with SR-OSPF adjacency SID tunnels.

These LFA policy attributes are checked against the outgoing interface used by the LFA/RLFA/TI-LFA backup path and selects the backup paths which satisfy the policy attributes.

The LFA policy indirectly applies to IPv4 LDP FECs when the LDP **fast-reroute backup-sr-tunnel** option is enabled, and the FEC is protected with an SR tunnel.

3.8.3.8 OSPF External Type 1 Overload

Release 19.7.R1 extends the support for OSPF overload condition to support External Type 1 routes. Prior to Release 19.7.R1, OSPF/OSPFv3 has supported overload conditions for internal routes and External Type 2 routes only.

3.8.3.9 RIB-API Egress Statistics

Release 19.7.R1 introduces the capability to collect, on egress data-path, statistics (in octets and packets) of traffic which is forwarded using RIB-API entries (more specifically, entries of the IPv4, IPv6, and MPLS tunnel tables). Statistics are collected without forwarding-class or QoS distinction and are available at the granularity of next-hops. An aggregate is also provided per instance of the entry. Statistics collection is configured as part of programming a RIB-API entry. These statistics can be displayed using **show** and **monitor** commands.

3.8.4 MPLS

3.8.4.1 Local CSPF Path Computation for SR-TE LSP

Release 19.7.R1 introduces full CSPF path computation for SR-TE LSP paths. This is an additional path computation method. Operators can now select among the hop-to-label translation, the local CSPF, or the PCE for a configured SR-TE LSP. The PCE option is not supported with the SR-TE LSP template.

The local CSPF feature supports the following capabilities:

- IPv4 SR-TE LSP and IPv4 SR-TE auto-LSP using SR-TE LSP template
- Path computation in single area OSPFv2 and IS-IS IGP instances
- Computation of full explicit TE paths using TE links as hops and returning a list of SIDs consisting of the adjacency SID and/or the parallel adjacency set SID
- Use random path selection in the presence of ECMP paths, satisfying the LSP and path constraints
- Provides an option to select protected or unprotected adjacency SIDs in the label stack of the computed path

- Provides an option to reduce or compress the label stack such that adjacency SIDs corresponding to a segment of the explicit path are replaced with a node SID whenever the constraints of the path are met by all the ECMP paths to that node SID
- Timer re-optimization of the active path of the SR-TE LSP, using the same mechanism applied to RSVP-TE LSP
- Manual re-optimization of a path of the SR-TE LSP, using the same mechanism applied to RSVP-TE LSP
- Support for unnumbered interfaces in the path computation
- Support for admin-group, hop-count, IGP metric, and TE-metric constraints

Note that several CLI commands for SR-TE LSP have been redesigned with the addition of newer commands which overlap with existing commands. The existing commands will be removed in a subsequent release. See Changed or Removed Commands in Classic Interfaces for more details.

3.8.4.2 MPLS Forwarding Policy Egress Statistics

Release 19.7.R1 introduces the capability to collect, on egress data-path, statistics (in octets and packets) of traffic which is forwarded using MPLS forwarding policies, more specifically both **binding-label** and **endpoint** policies the next-hops of which have been configured with a **resolution-type direct**. Statistics are collected without forwarding-class or QoS distinction and are available at the granularity of next-hops. An aggregate is also provided per instance of the policy. Statistics collection is configured using the **config**>**router**>**mpls**>**fwd-policies**>**fwd-policy**>**egress-statistics** command. These statistics can be displayed using **show** and **monitor** commands.

3.8.4.3 S-BFD for SR-TE LSPs

Release 19.7.R1 introduces Seamless BFD (S-BFD) for SR-TE LSPs. Unlike LSP BFD, S-BFD does not rely on the traditional BFD session bootstrapping process (handshake) or session state at the tail end of a session. Instead, when S-BFD is initialized, a set of discriminators are selected by the system for specific purposes (reflector or initiator).

The S-BFD reflector is configured once per system under the new **config**>**bfd**>**seamless-bfd** context. A mapping between reflector discriminators and their IP address is configured under the new **config**>**router**>**bfd**>**seamless-bfd** context, and configuration of the head end of the session under the **config**>**router**>**mpls**>**lsp**>**bfd**, **config**>**router**>**mpls**>**lsp**>**primary**>**bfd**, and **config**>**router**>**mpls**>**lsp**>**secondary**>**bfd** contexts.

S-BFD provides a mechanism to check the data path forwarding for an SR-TE LSP. The system supports a number of consequent actions if the S-BFD session fails, which are configured using the **config**>**router**>**mpls**>**lsp**>**bfd**>**failure-action** command. The SR-TE LSP may be configured to switch to the next available standby or secondary path, take the LSP operationally down using the new **failover-or-down** option, or to simply raise a trap if the new none option is used.

S-BFD control packet timers may be configured down to 10ms.

3.8.4.4 Secondary Path for SR-TE LSPs

Release 19.7.R1 adds the support for secondary and standby paths to SR-TE LSPs. These paths are used to provide end-to-end protection for the SR-TE LSP. Standby paths are permanently programmed in the data plane, while non-standby secondary paths are only programmed when needed. Up to a maximum of three paths per LSP are supported. Secondary and standby paths are configured using the **config**>**router**>**mpls**>**lsp** *lsp-name* **sr-te**>**secondary** context. This feature is supported with both the hop-to-label translation and the local CSPF path computation methods.

See Limited Support Features for more information about the local CSPF support.

3.8.4.5 Weighted ECMP Enhancements for LDP-over-RSVP

Release 19.7.R1 adds the support for weighted ECMP for LDP-over-RSVP, where the LDP FEC is resolved to a static-route that also resolves to one or more RSVP tunnels. The next-hops for the static-route may be direct or indirect, and in the indirect case, may be the same or different. The support is also added for weighted ECMP in classic LDP-over-RSVP where the **prefer-tunnel-in-tunnel** option is configured for LDP.

Release 19.7.R1 also increases the granularity of the normalization of ECMP weights used for LDP-over-RSVP from 32 to 64 in both instances, and in the existing implementation of weighted ECMP for LDP-over-RSVP where the LDP FEC resolves to an IGP shortcut tunnel.

3.8.5 Services

3.8.5.1 EVPN-VXLAN Multi-homing Along with Non-system IPv4/IPv6 Termination

Release 19.7.R1 adds the support for EVPN single-active and all-active multi-homing in VXLAN VPLS and Routed-VPLS (R-VPLS) services when VXLAN tunnels are terminated on non-system IPv4 and IPv6 addresses (VTEPs).

3.8.5.2 EVPN-VXLAN Multi-homing Support on R-VPLS

Release 19.7.R1 adds the support for EVPN single-active and all-active multi-homing in VXLAN Routed-VPLS (R-VPLS) services with a single BGP-EVPN instance. Port, LAG, and SDP-based Ethernet Segments (ES) and virtual Ethernet Segments (vES) are supported.

3.8.6 Subscriber Management

3.8.6.1 Dynamic Policer Parameter Overrides for PCC Rules

Release 19.7.R1 adds the support for per-PCC rule overrides of the following dynamic policer parameters: **parent** *arbiter-name* [**weight** *weight-level*] [**level** *level*], **mbs**, **cbs**, **stat-mode**, and **packet-byte-offset**. The overrides can be specified at the activation of a PCC-rule-based subscriber service or when installing a flow-based PCC rule using Diameter Gx.

3.8.6.2 Priority-based Server Selection for ISA RADIUS Policies

Release 19.7.R1 introduces a **direct-priority** mode server selection to ISA RADIUS policies. In this mode, the ISA always sends requests to the highest-priority server available.

3.8.6.3 NGE Support for WLAN-GW L2oMPLSoGRE

Release 19.7.R1 introduces NGE-encrypted services for WLAN APs that support NGE and L2oMPLSoGRE transport of WLAN traffic to the WLAN-GW.

3.8.7 OAM

3.8.7.1 OAM-PM Delay Metrics Average Streaming

Release 19.7.R1 introduces the support for OAM-PM to configure delay-templates for streaming average Frame Delay (FD) and Inter-Frame Delay Variation (IFDV) computed over shorter durations of time than is typical for Service Level Agreement (SLA) measurements. SLAs are typically measured using measurement-intervals of five (5) minutes or longer. Streamed results for network optimization rely on shorter samples windows. The delay-templates include various parameters for common configuration elements including the metric type and direction, the length of the sample window, and window integrity percentage. This template can be associated with the various OAM-PM delay tests, for Ethernet (DMM), IP (TWAMP-Light), and MPLS (DM).

3.8.8 VSR

3.8.8.1 Filter Rate-Limit Action

Release 19.7.R1 introduces the support for IPv4, IPv6 filter **action rate-limit** using ingress or egress filter policies on 7750 VSR.

3.8.8.2 IPsec Secure Interface

Release 19.7.R1 introduces a secured interface to allow users to create one or multiple static IPsec tunnels under an IP interface to be used for secure traffic forwarded via the interface. The local IPsec tunnel endpoint address is one of interface addresses. By default, all ingress traffic (except for MPLS/SDP/SSH traffic) are subject to IPsec processing. However, users can optionally use an **ip-exception** filter to specify the ingress traffic and bypass the IPsec processing.

This feature is only available on the VSR-I platform.

3.8.8.3 vSeGW External Offload

In Release 19.7.R1, VSR-I supports IPsec fastpath offloading using Intel QuickAssist (QAT) technology. The system automatically detects and utilizes QAT hardware when it is made available to the VSR. When enabled, ESP packet encryption and decryption are processed by QAT hardware. This further improves IPsec throughput on VSR.

This feature supports the following QAT hardware on KVM hypervisor via SR-IOV:

- Intel PCH chipset C627/C628
- Intel QuickAssist Adapter 8970

3.9 Release 19.5.R2

There are no new major features in Release 19.5.R2.

3.10 Release 19.5.R1

The following sections describe the new features added in Release 19.5.R1 of SR OS.

- Hardware
- System
- Routing
- MPLS

- Services
- Subscriber Management
- Application Assurance

3.10.1 Hardware

3.10.1.1 Power-save Mode

Release 19.5.R1 introduces the ability to place an IOM/IMM/XCM into **power-save** mode when not in use. In **power-save** mode, the IOM/IMM/XCM consumes minimal power while installed and provisioned in a system.

Individual MDAs/XMAs can not be put in **power-save** mode and will not be functional when installed in an IOM/XCM in **power-save**.

Power-save mode can be configured using the **config**>**card**>**power-save** command. The card placed in **power-save** mode is forced into an idle state to consume minimal power. To bring up a card, use the **no power-save** command to allow the card to power up and initialize. By default, all slots and cards are set to **no power-save**.

3.10.1.2 Support for 7210 SAS-Sx/S/MxP Ethernet Satellites Running Release 11.0 Software

A 7750 SR or 7950 XRS chassis running SR OS Release 19.5.R1 can act as the host to a 7210 SAS-Sx/SAS-S/SAS-MxP running 7210 SAS OS Release 11.0 and running in satellite mode. Ethernet satellites running 7210 SAS OS Release 11.0 support the same functionality as with Release 10.0. Chassis running SR OS Release 19.5.R1 can also act as hosts to 7210 Ethernet satellites running Releases 9.0 and 10.0.

3.10.1.3 7750 SR 3-port CFP2-DCO MDA-e-XP

Release 19.5.R1 introduces the 7750 SR 3-port CFP2-DCO MDA-e-XP. This MDA-e-XP is supported on the 7750 SR-12e (with all variants of the IOM5-e) and 7750 SR-1 chassis and supports a wide range of pluggable interfaces. The 7750 SR 3-port CFP2-DCO MDA-e-XP supports the following modes:

- Long Haul
- Metro
- Interop
- Interop3

The MDA-e-XP is not supported on the 7750 SR-7-B/12-B chassis. Note that the CFP2 100G Base-SR10 optical module is not supported on this 7750 SR 3-port CFP2-DCO MDA-e-XP.

3.10.1.4 7750 SR 10-port 10/1GE MACsec MDA-a-XP

Release 19.5.R1 introduces the 10-port 10/1GE MACsec MDA-a-XP for the 7750 SR-a4/a8 chassis. This MDA-a-XP supports a wide range of pluggable interfaces. The key feature for this MDA-a-XP is MACsec.

The 10/100/1000-T SFPs (3HE00062CB or 3HE11904AA) are not supported in this MDA-a-XP.

3.10.2 System

3.10.2.1 7450 ESS Mixed-mode Chassis Behavior Default

In Release 19.5.R1, the 7450 ESS mixed-mode becomes the default behavior for all 7450 ESS chassis. This mode of operation allows a 7450 ESS chassis to support all 7750 SR functionality. As of Release 19.5.R1, all supported IOMs and IMMs for the 7450 ESS chassis are 7750 SR-capable, so the 7450 ESS chassis will always operate in this mode.

The CLI commands previously used to configure this feature have been removed from the SR OS. No administrator action is necessary as these commands will be automatically removed from the operating configuration during the upgrade to Release 19.5.R1. In addition, when the configuration is saved under Release 19.5.R1 or higher, none of these commands are included in the save configuration file. The associated SNMP MIB entities associated with these commands have also been obsoleted.

If 7750 SR features are to be used on a 7450 ESS, the appropriate use licenses must be purchased. Please contact your Nokia representative to assist in this process.

3.10.2.2 Cflowd: Variable Sampling per Interface

Release 19.5.R1 enhances the configuration of Cflowd sampling rates to allow for the creation of up to five sampling profiles, each of which can specify a different Cflowd sampling rate. The sample profiles can then be associated with individual IP interfaces to control the rate at which traffic is sampled on that particular interface.

The new sample profiles are created using the **config**>**cflowd**>**sample-profile** *profile* [**create**]>**sample-rate** *rate* CLI command.

A **sample-profile 1** is automatically created and has a default *rate* of 1000. If the system Cflowd sampling rate had been changed to a non-default value, then that value is translated to the new **sample-profile 1 sample-rate** *rate*. Only one of the configured sample profiles can have a sample rate of 256 or less.

A sample profile is then associated with the IP interface in a new optional parameter in the sampling commands under IP interfaces in the **config**>**router**>**interface**>**cflowd-parameters**>**sampling** context.

3.10.2.3 Distributed CPU Protection for R-VPLS

Release 19.5.R1 introduces broadcast and multicast traffic rate limiting using Distributed CPU Protection (DCP) in Routed-VPLS (R-VPLS) services. In prior releases, broadcast and multicast traffic was not subject to DCP in R-VPLS services.

3.10.2.4 Encryption Method for Imported Certificate Key

Release 19.5.R1 introduces an improved secure format for imported certificate, key, and CRL files in a CF3:/system-pki directory. Newly imported certificates, keys, or CRLs use this new format, while the system continues to load existing imported files in the legacy format.

This behavior can be changed using the **config**>**system**>**security**>**pki**>**imported-format secure** command which enforces the system to only load imported files in the new format.

The new **admin**>**certificate**>**convert-file** command converts imported files between the legacy format and new format.

With the new secure format:

- a stronger encryption algorithm is utilized
- imported certificates and keys both are now encrypted
- the internal key for encryption is now chassis-specific
- a compressed format is now used for imported CRL files to save space

3.10.2.5 Filter: Match Packet Length

Release 19.5.R1 introduces the capability to match on the total IPv4 or IPv6 packet-length using ingress and egress IPv4 or IPv6 filter policy. This new match criterion is configurable using the filter policy type **packet-length**.

3.10.2.6 Filter: IPv4 and IPv6 SRC-MAC Filtering in VPLS

Release 19.5.R1 introduces the capability to match on the source MAC address using egress IPv4 or IPv6 filter policy in VPLS services. This new match criterion is configurable using the filter policy type **src-mac**.

3.10.2.7 Filter: Pattern-based Whitelisting

Release 19.5.R1 introduces support for pattern-based whitelisting using ingress IPv4 and IPv6 filtering. This capability is enabled using the **forward-when pattern** filter action. The pattern match is defined by a hexadecimal **expression**, **mask**, **offset-type**, and **offset-value**.

3.10.2.8 LAG Hash Weight

Release 19.5.R1 introduces LAG port **hash-weight** to customize the flow distribution hashing between LAG ports. The LAG port **hash-weight** value is normalized internally to distribute flows between LAG ports if all ports have a **hash-weight** configured.

3.10.2.9 100GE/400GE Mixed-Speed LAG

Release 19.5.R1 introduces the support for mixed-speed LAG for 100GE and 400GE ports. This capability does not require **port-weight-speed** and is compatible with LAG port **hash-weight.**

3.10.2.10 Node Management (NETCONF and gRPC) in VRF/VPRN

Release 19.5.R1 introduces the ability to access an SR OS router via NETCONF/gRPC using a VRF/VPRN. This is achieved by configuring SR OS to allow NETCONF/gRPC to use a management VRF/VPRN, therefore allowing SR OS NETCONF/gRPC to become VRF-aware.

3.10.2.11 Node Discovery using OSPF

Release 19.5.R1 introduces the ability to discover an SR OS node using OSPF and type 10 Opaque LSA. This node discovery is available for VPRN only. The node is configured using a network element profile, which has all necessary information for node discovery, including node NEID, NEIP, name, chassis type, and MAC address. The network element profile can be added to an OSPF area. Once added, the network element information is advertised using OSPF type 10 Opaque LSA to the rest of the nodes. An aggregation node (the node closest to the Network Management System (NMS)) gathers all of the network element information received from OSPF and converts it into MIB information. The aggregation node also generates traps as necessary to update the NMS. In addition, the NMS can walk the MIB table to update its new view of the network.

3.10.2.12 Local Boot for 7950 XRS-40 Extension Chassis CPM

In Release 19.5.R1, the boot and upgrade process of the extension chassis in the 7950 XRS-40 has changed to use local versions of the boot loader, BOF (Boot Option File), and image files. The image files must be copied onto the local cf3: of the extension chassis before an upgrade. This copy is included in the execution of the **admin redundancy synchronize boot-env** command of Release 15.0.R1 and higher.

3.10.3 Routing

3.10.3.1 Default EBGP Route Propagation Behavior without Policies

Release 19.5.R1 introduces the support for default external BGP (EBGP) route propagation behavior without policies as described in RFC 8212.

The behavior of a newly-created or existing routing instance, group, or EBGP neighbor in a classic interface (the classic CLI and SNMP) without import or export policies maintains backwards compatibility with the insecure default to advertise and receive all routes. It is not compliant with RFC 8212. The secure default behavior must be enabled using the **ebgp-default-reject-policy** command in these cases.

The behavior of a newly created routing instance, group, or EBGP neighbor via a model-driven interface (the MD-CLI, NETCONF, or gRPC) without import or export policies applies the secure default behavior to reject all routes. It is compliant with RFC 8212. The secure behavior can be disabled using the **ebgp-default-reject-policy** command. However, Nokia recommends configuring import and export policies that express the intended routing instead of using the insecure default behavior.

3.10.3.2 IS-IS and OSPF Reference Bandwidth Increase

Release 19.5.R1 increases the IS-IS and OSPF reference-bandwidth from 4,294,967,295 (equal to the highest 32-bit unsigned integer value) to 18,446,744,073,709,551,615 (equal to the highest 64-bit unsigned integer value). This increased value range allows meaningful derived auto-cost for higher speeds (for example, a 64 x 100GE LAG).

3.10.3.3 Multiple Peer AS Support for Dynamic BGP Sessions

Release 19.5.R1 allows dynamic BGP sessions associated with a single BGP peer-group to belong to different peer AS (autonomous systems). In previous SR OS releases, those dynamic neighbors were required to have the same peer AS. In Release 19.5.R1, a dynamic BGP session can be rejected if the neighbor does not report an AS number in an allowed list. The allowed ranges are configured against the IP prefix used to match the neighbor IP address.

3.10.3.4 Node Protection with Remote LFA (with LDP-SR Stitching Support)

In Release 19.5.R1, the node protection option of the Remote LFA (RLFA) and Topology-Independent LFA (TI-LFA) features are now supported with an LDP FEC that uses an SR tunnel as a LFA backup. This applies when the **ldp fast-reroute backup-sr-tunnel** option is enabled and the SR tunnel used to protect the LDP FEC is itself protected with RLFA or TI-LFA node-protection backup.

3.10.3.5 RFC 5549 Support for IPv4 BGP Routes

Release 19.5.R1 introduces the ability for the SR OS router to advertise and receive BGP routes that convey reachability to IPv4-unicast destinations but are reachable through IPv6 next-hops. This capability is allowed by the extensions in RFC 5549 and is useful in networks or regions with IPv6-only interfaces.

3.10.3.6 Control Over the BGP Multi-path Algorithm

Release 19.5.R1 adds additional controls over the BGP multi-path algorithm. Enhancements include:

- The ability to specify multi-path parameters on a per-address-family basis (IPv4, label-IPv4, IPv6, and label-IPv6 routes).
- Support for an **unequal-cost** option. When this option is specified there is no change to the best-path selection criteria (routes are still ordered by the next-hop-cost when the higher criteria are equal). When multi-paths are selected, BGP ignores the differences in the next-hop cost.
- Support for marking peers and groups as **multipath-eligible**. This implements support for selective multi-path. Selective multi-path restricts the set of peers that can provide ECMP paths for IP prefixes. It is supported for IPv4, label-IPv4, IPv6 and label-IPv6 routes.

3.10.4 MPLS

3.10.4.1 Class-based Forwarding and ECMP for LDP FECs Resolving to IGP Shortcuts

Release 19.5.R1 introduces the capability to perform both Class-based Forwarding (CBF) and ECMP for LDP FECs resolved to RSVP-TE LSPs as IGP shortcuts at an LDP LSR. Previously, it was only possible to perform CBF (that is, no ECMP). Two methods exist to configure LSPs with CBF information. The direct FC-to-LSP configuration does not inherit ECMP capabilities; however, the FC-to-Set configuration does inherit the ECMP capability, and the system is subject to perform ECMP (expected behavior) per forwarding set when the system is upgraded to Release 19.5.R1.

3.10.4.2 Intra-AS Option B Support for MVPN

Release 19.5.R1 introduces the **configure router ldp import-pmsi-routes mvpn** and **mvpn-no-export-community** CLI commands to control the caching of the BGP routes.

The commands are disabled by default. However, **mvpn** is automatically enabled upon upgrade to Release 19.5.R1 so MVPN inter-AS functionality from prior releases is still compatible with Release 19.5.R1. On a new Release 19.5.R1 SR OS node, the **mvpn** command must be executed to enable the inter-as option B functionality for MVPN.

3.10.5 Services

3.10.5.1 Affect LAG Signaling Based on EVPN-VPWS Oper-group Events

Release 19.5.R1 introduces the support of **monitor oper-group**s in LAGs and **epipe**>**bgp-evpn**>**mpls**>**oper-group** *name*. They can be used together in the two PEs attached to an EVPN-VPWS service to provide a link loss forwarding functionality. The withdrawal of a remote Auto-Discovery per-EVI route can bring down the BGP-EVPN **oper-group**, and therefore the LAG that is monitoring the **oper-group** can trigger the LAG standby signaling toward the connected CE, propagating a remote failure to the local CE. The LAG signaling can be based on **lacp** or **power-off**.

3.10.5.2 EVPN-VXLAN Multihoming

Release 19.5.R1 introduces the support for EVPN single-active and all-active multi-homing in VXLAN VPLS services with a single BGP-EVPN instance. Port, LAG and SDP-based Ethernet-Segments (ES) and virtual Ethernet Segments (vES) are now supported on VXLAN services. The following considerations apply:

- All existing Designated Forwarder (DF) election algorithms and options are supported for VXLAN services on PEs attached to an ES.
- The multi-homing supported procedures are compliant with RFC 8365. In particular, forwarding and Split Horizon filtering are based on the “Local Bias” procedures specified in RFC 8365.
- In addition, aliasing and backup functions to remote PEs that are attached to an ES are now supported.

- The multi-homing capabilities are enabled in all the PEs attached to the VPLS service by configuring the commands **vpls**>**bgp-evpn**>**vxlan**>**auto-disc-route-advertisement** and **vpls**>**bgp-evpn**>**vxlan**>**mh-mode network**. These two commands enable the advertising and processing of multi-homing routes (the former command) and activate the DF election procedures (the latter command).

See Known Limitations for more information.

3.10.5.3 IPv4 MPLS Forwarding Policy and IPv4 RIB API Tunnel Resolution for EVPN MPLS Services

Release 19.5.R1 introduces the support for two new IPv4-based tunnel types for EVPN MPLS auto-bind resolution:

- IPv4 MPLS forwarding policy
- IPv4 RIB API

These two tunnel types are added to the auto-bind resolution and any auto-bind resolution filter support. The new tunnel types behave as any other MPLS tunnel type as far as EVPN is concerned. In particular:

- MPLS forwarding policy and RIB API tunnels are supported in all the EVPN services that support MPLS tunnels (such as VPLS, Epipe, R-VPLS, B-VPLS, and dual BGP instance services).
- There are no restrictions in terms of features supported in EVPN-MPLS services due to the new supported tunnels.
- EVPN multi-homing is supported irrespective of the MPLS tunnels selected in the auto-bind resolution process and including MPLS forwarding policy and RIB API tunnels.

3.10.5.4 IGMP/PIM Snooping and MCS Support along with EVPN-VXLAN Multi-homing in VPLS

Release 19.5.R1 introduces support for IGMP and PIM snooping on VPLS services that are configured to support EVPN-VXLAN multi-homing. In addition, IGMP and PIM (for IPv4) snooping states can be synchronized on the Ethernet-Segment peer PEs by using Multi-Chassis Synchronization (MCS).

3.10.5.5 Video Application: Perfect Stream Support

Release 19.5.R1 introduces the support for Perfect Stream (also known as dual stream selection) on MS-ISA2 in addition to MS-ISA. Perfect Stream utilizes two diverse network path multicast streams to generate a "perfect" multicast stream, mitigating packet losses induced by the network. Perfect Stream can be configured together with VQM on the same ISA to monitor and compare the incoming impaired streams, as well as the repaired stream.

3.10.6 Subscriber Management

3.10.6.1 WLAN-GW UE and Tunnel Query Persistency

Release 19.5.R1 introduces full CLI and resiliency support to the previously SNMP-only query-based UE and tunnel state retrieval. Queries can be configured under **configure subscriber-mgmt wlan-gw ue-query** and **tunnel-query**; results can be displayed via the **show subscriber-mgmt wlan-gw ue query-results** and **tunnels query-results** commands.

3.10.6.2 Subscriber Management Support over Redundant Uplinks on Ethernet Satellites

Release 19.5.R1 adds the support dual Ethernet satellite nodes with redundant uplinks in subscriber management.

3.10.6.3 Per-host DNS Override Using Destination NAT

Release 19.5.R1 introduces the ability to enable and disable the overriding of the DNS address on a per-host basis. The DNS address override is enabled or disabled using a RADIUS VSA in an Access-Accept or CoA message corresponding to a host. This allows the capability to force DNS packets of a device in a Virtual Residential Gateway (vRGW) to the DNS servers of its choice. DNS override is achieved by subjecting DNS traffic to destination NAT (DNAT). The traffic that is subjected to

DNAT is selected by applying a NAT classifier. Also, support is added for RADIUS VSA to specify a per-host default address for overriding the address in DNS packets. Additionally, the feature extends existing NAT classifier match criteria to also include the destination IP address in selecting traffic that goes through DNAT for DNS override.

3.10.7 Application Assurance

3.10.7.1 AA vRGW Nested Router Detection

Release 19.5.R1 introduces the ability to detect nested routers in a Bridged Residential Gateway (BRG) home. A nested router is detected using multi-device-detection for an AA-subscriber device in the home. The implementation for this scenario uses the AA capability for tethering detection on an ESM-MAC AA-subscriber.

4 Enhancements

The following sections describe new enhancements in SR OS releases. Enhancements from Releases 16.0.R2 to 16.0.R7 also apply to Release 19.*x*. Refer to the most recent *SR OS 16.0 Release Notes* for the summary of enhancements in Releases 16.0.R2 through 16.0.R7.

Note:

- For the list of new and updated Application Assurance protocols and applications supported in Release 19.10.R6, see the following spreadsheet at the Nokia online customer support site:

  SR OS 19.x 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 19.10.R6

4.1.1 Routing

- Release 19.10.R6 enhances IS-IS to keep preferring local RSVP tunnels over local IGP routes when the node goes into an overload state. [329916]

4.1.2 MPLS

- Release 19.10.R6 adds the ability to configure a timeout action for LSP self ping using the new **timeout-action retry|switch** command under the **config**>**router**>**mpls**>**lsp-self-ping** context. When an LSP self ping session times out, the router either retries the path (default behavior) or switches to the new path. [339509]

4.1.3 Enhancements Added in Prior Releases

See the items marked **[NEW]** in Release 19.10.R1 and Release 19.10.R3.

4.2 Release 19.10.R5

There are no new enhancements in Release 19.10.R5.

4.3 Release 19.10.R4

There are no new enhancements in Release 19.10.R4.

4.4 Release 19.10.R3

4.4.1 System

- Release 19.10.R3 introduces a new **tools perform chassis link-check** command that verifies the connection between the XMA and its respective XCM on the 7750 SR-2s, SR-7s, and SR-14s platforms. [331984]
- Release 19.10.R3 introduces a new compatibility mode of **long-haul-non-diff** for the CFP2-DCO optical module. This is equivalent to the long-haul mode but uses non-differential encoding of the data. In addition, the minimum **rx-los-threshold** has been lowered to -30 dBm for these modules. [336816]

4.4.2 Configuration Coverage in Model-driven Interfaces

This section lists configuration commands that were unavailable in MD interfaces in previous releases, but are newly available for configuration in Release 19.10.R3. These are features and configuration commands that existed in classic CLI and SNMP in previous releases.

Refer to Unsupported Configuration in MD Interfaces for more details about configuration elements that are unsupported.

4.4.2.1 System

- MACsec

4.4.3 Routing

- IPv6 Traffic Engineering (IS-IS) is now ready for production networks. This feature was introduced in Release 19.10.R1.
- Release 19.10.R3 extends the LFA policy feature to IPv4 and IPv6 SR-ISIS adjacency SID tunnels. **[NEW]**

4.4.4 MPLS

- IPv6 SR-TE LSP is now ready for production networks. This feature was introduced in Release 19.10.R1.

4.4.5 Services

- EVPN-MPLS Layer-3 OISM is now ready for production networks. This feature was introduced in Release 19.10.R1.
- IPv6 Tunnel Resolution for EVPN MPLS Services using SR-TE tunnels is now ready for production networks. This enhancement was introduced in Release 19.10.R1.
- EVPN IPv4 Host Mobility within the Same R-VPLS Service is now ready for production networks. This feature was introduced in Release 19.10.R1.
- Release 19.10.R3 adds an object to the sapPortStateChangeProcessed SNMP trap to indicate if the change is due to a port Up/Down/Other event. [336822]

4.4.6 Subscriber Management

- Multi-Chassis Synchronization Support for Usage-Monitoring (Gx) is now ready for production networks. This feature was introduced in Release 19.10.R1.

4.4.7 Application Assurance

- Release 19.10.R3 adds a new show command for traffic-type tls-family for statistics to track the number of sessions, packets and bytes using TLS1.3 with encrypted SNI. [332763]
- In Release 19.10.R3, support is provided for an ESNI flow-attribute that provides reporting and policy control for flows that use an encrypted server name indication (eSNI). [336125]

4.4.8 NAT

- Release 19.10.R3 enhances NAT PCP where a NAT subscriber is forced to use a single external IP address for a given NAT policy, regardless of the suggested (requested) port in a PCP MAP request.

  The new behavior must be explicitly enabled with the **config**>**service**>**nat**>**pcp-server-policy**>**reuse-ext-ip** command. [330966]

4.5 Release 19.10.R2

4.5.1 MD-CLI

- The MD-CLI **admin system license** commands are now ready for production networks. This enhancement was introduced in Release 19.10.R1.

4.5.2 Routing

- BGP Selective Label-IPv4 Route Installation is now ready for production networks. This feature was introduced in Release 19.10.R1.

- Micro-Loop Avoidance Using Loop-free SR Tunnels (IS-IS) is now ready for production networks. This feature was introduced in Release 19.10.R1.

4.5.3 Services

- EVPN-VXLAN Network-interconnect Multi-homing (in Two VXLAN Instance Services) is now ready for production networks. This feature was introduced in Release 19.10.R1.
- EVPN Multi-Homing Support for MPLS Tunnels using IPv6 or Non-system IPv4 Addresses for the following tunnel types is now ready for production networks. This feature was introduced in Release 19.10.R1.
  - mpls-fwd-policies
  - ldp
  - rib-api

4.5.4 QoS

- Release 19.10.R2 adds a new port scheduler H-QoS algorithm which is supported for use on both Ethernet Vports and Ethernet physical ports. The new algorithm can be enabled within the port scheduler policy configuration by setting the **hqos-algorithm** to **above-offered-allowance-control**. When the new algorithm is enabled, parenting queues and schedulers to the port scheduler is supported. The new algorithm allows control of the amount of bandwidth in excess of the offered rate that is given to a queue. Tuning of the algorithm is achieved using new commands under the **above-offered-allowance** context in the advanced configuration policy that is applied to a queue.

4.5.5 Subscriber Management

- ESM Subscriber Unicast Reserve Bandwidth is now ready for production networks. This enhancement was introduced in Release 19.10.R1.

4.6 Release 19.10.R1

4.6.1 Hardware

- Release 19.10.R1 enables the support for faster HBMv2 memories. This enhancement improves performance on the 7950 XRS XMA - 7950 XRS 2.4T 12pt QSFP-DD and XMA - 7950 XRS 6pt CFP2 DCO as of this release. Support for faster HBMv2 memories enables:
  - Enhanced PPS rates on all platforms and line cards when shipped with faster HBMv2
  - Throughput enhancements on select platforms

  If additional details are required regarding specific platforms or line cards post Release 19.10.R1, contact Nokia support.
- Release 19.10.R1 introduces a new **internal-frame-loss** event type that can be used under the **card**>**mda**>**event** context. First, an **event** must be created and the event type can be specified as **internal-frame-loss**. This new event provides options on how to handle internal frame loss events on FP2- and FP3-based MDAs/XMAs. The **event** can have the following actions to be taken on the MDA/XMA when a internal frame loss event occurs. The actions are:
  - **no action** – Do nothing (basic error counting)
  - **log-only** – Generates a log-event, SNMP trap
  - **reset-mda** – Resets to clear the error
  - **fail-mda** – Fails the MDA/XMA until operator intervention

  The default action will be **reset-mda** in an attempt to clear any internal conditions that are causing the frame loss. If the **event** is not created, no action will be taken. [240787]
- Release 19.10.R1 adds the Soft Reset support for the 7750 SR 12-port 10/1GE MACsec SFP+ MDA-e (me12-10/1gb-sfp+). [268287]
- Release 19.10.R1 adds the Soft Reset support for the 7750 SR 10-port 10/1GE MACsec MDA-a-XP (maxp10-10/1gb-msec-sfp+). [329595]
- Release 19.10.R1 adds the Soft Reset support for the following FP4-based hardware assemblies. This provides the foundation for ISSU support on FP4-based systems (7750 SR-2s/7s/14s) with Release 19.10.R1 as the source/starting point of the ISSU.
  - All XCM for the 7750 SR-2s/7s/14s:
    - xcm-2s
    - xcm-7s

        - xcm-14s
    - The following XMAs for the 7750 SR-2s/7s/14s:
        - s18-100gb-qsfp28
        - s36-100gb-qsfp28
        - s36-400gb-qsfpdd
    - XCM-2 for the 7950 XRS-20 and 20e (xcm2-x20)
    - The following XMA for the 7950 XRS:
        - x24-100g-qsfp28
        - x12-400g-qsfpdd
        - x6-200g-cfp2-dco
        - x6-400g-cfp8
    - IOM5-e for the 7750 SR-7/12/12e (iom5-e)
    - The following MDA-e-XP for the 7750 SR:
        - me6-100gb-qsfp28
        - me12-100gb-qsfp28
        - me3-200gb-cfp2-dco

4.6.2 Satellites

- In Release 19.10.R1, the configurable range for the **config**>**system**>**satellite**>**eth-sat**>**client down-delay** <*seconds*> command has been changed to allow values of 0-1800 seconds, instead of the original 5-1800 seconds. [326754]

4.6.3 System

- Release 19.10.R1 introduces a new tools command, **tools**>**dump**>**filter**>**cam-utilization** [**card** *slot-number*], to display the ACL filter resource utilization per line card and per FP CAM on FP4-based systems. The system also raises a trap when a specified FP4 CAM utilization goes beyond 85%.

  In Release 19.10.R1, after the first overload condition is detected for a specified ACL filter FP4 CAM, the system now interactively rejects the addition of entries or policies on the same FP CAM.
- Release 19.10.R1 introduces the following gRPC command accounting log events:

  - grpc_auth
  - grpc_unauth

  One of these log events is sent upon receiving any RPC from a gRPC client. Both events are mapped to the sros-md-rpc-accounting-event NETCONF notification.

- Release 19.10.R1 introduces “Bytes” encoding in telemetry notifications.
- Release 19.10.R1 includes model-driven **system thresholds** state information for dynamic threshold objects (**cflash-cap-alarm-percent**, **cflash-cap-warn-percent**, **kb-memory-use-alarm**, **kb-memory-use-warn**). [323255]
- In Release 19.10.R1, by enabling **config**>**log**>**services-all-events service** *service-id*, all log events will be enabled and sent to syslog via the corresponding service (VPRN). Prior to Release 19.10.R1, the syslog could be enabled in the VPRN using the **config**>**service**>**vprn**>**log**>**syslog** command, but events were limited to a subset of VPRN events and would not send all events to the syslog. This included the security and debug events. [326537]
- Release 19.10.R1 introduces the ability for a triggered EHS/CRON script to execute while there is a datastore lock (that is, started by a MD interface) in place. The **configure system script-control script-policy** *policy-name* **lock-override** command is introduced in the classic CLI and the MD-CLI to control if a triggered EHS/CRON script should bypass an existing MD interface’s datastore lock or not.

4.6.4 MD-CLI

- Release 19.10.R1 introduces the ability to display the present working context (**pwc**) in alternative display formats. These formats include a format similar to XPath that can be utilized with telemetry systems and a YANG modeled path format that can be utilized with RESTCONF-based management systems.
- Release 19.10.R1 introduces command extensions to allow the configuration to be displayed in JSON format. The **admin show configuration json** command shows the running configuration in JSON format, while the extensions to the **info** command, used within configuration mode, shows the configuration from a specific working context in JSON format.

  The following info extensions are added: **info json**, **info inheritance json**, and **info converted json**.
- Release 19.10.R1 adds the ability to insert entries into a policy statement or OpenConfig routing policy. Entries can be inserted into these user-ordered lists either before or after a specific entry or at the beginning or end of the list.

To use this feature, the list must already exist. The **insert** command can also be used to move an already existing item within the same list.

- Release 19.10.R1 introduces the **traceroute** command in the MD-CLI to display the route that packets take to a host.
- Release 19.10.R1 introduces the **ping** command in the MD-CLI to verify the reachability of a host.
- In Release 19.10.R1, the MD-CLI **admin** command tree is extended to include the following commands and parameters to match the classic CLI functionality:
  - **clear**
  - **disconnect**
  - **reboot**
  - **redundancy**
  - **save**
  - **set**
  - **show**
  - **support-mode**
  - **system**

  See Limited Support Features for more information.
- In Release 19.10.R1, the command completion in MD-CLI keys to references is enhanced to display wildcard references. For example, interface names from the Base router and IES services (the wildcard reference) will be displayed in command completion under **configure router isis interface**. [320443]
- In Release 19.10.R1, the output of the **tree** command is enhanced to show unnamed parameters consistently with how unnamed keys are shown, so that the entire command syntax is displayed. [325052]
- In Release 19.10.R1, the MD-CLI command completion is enhanced to complete Boolean (true/false), enumeration, and **admin-state** (enable/disable) parameters with the Spacebar and Enter keys in addition to the Tab key. [325671]
- In Release 19.10.R1, a new configuration parameter **configure system management-interface cli md-cli command-accounting-during-load** is added to disable remote command accounting after issuing a load or rollback command, making the operation faster with a large configuration. [328956]
- In Release 19.10.R1, the MD-CLI command completion is extended to show keys in the same order that is displayed in the configuration, so that variable parameters area easier to see. Previous releases displayed the output in alphabetical order. [276758]

4.6.5 NETCONF

- Release 19.10.R1 enhances SR OS NETCONF support for <copy-config> by allowing the following scenarios:
  - <source>=<config> and <target>=<startup>
  - <source>=<config> and <target>=<url>
  - <source>=<statrup> and <target>=<candidate>
  - <source>=<url> and <target>=<candidate>
  - <source>=<url> and <target>=<url> (as long as both are not remote URLs)
  - <source>=<candidate> and <target>=<startup>
  - <source>=<candidate> and <target>=<url>
- In Release 19.10.R1, the **config**>**system**>**netconf**>**port** command can be set to 22 to allow NETCONF management in SR OS using port 22. The command applies for both VPRN/VRF NETCONF management and non-VPRN/VRF NETCONF management in SR OS. The default port to use for NETCONF management is port 830. Only one port can be configured/used for NETCONF management at any time.
- In Release 19.10.R1, when using the <validate> or <commit> RPCs, the SR OS NETCONF server will return multiple errors in the <rpc-error> if more than one error exists. This is an enhancement to the old behavior of returning one error at a time in the <rpc-error>. The old behavior forced users to fix the first found error, then re-send the <validate>/<commit> RPC before finding the next error and repeating the fix-error and re-send RPC procedure until no errors are returned.
- Release 19.10.R1 introduces the ability to use an <edit-config> RPC on a user-ordered list to control where an entry can be inserted in the list. The “insert” and “key” attributes are supported to insert and move a user-ordered list’s entry in the candidate datastore prior to a commit.
- Release 19.10.R1 introduces the following two NETCONF command accounting log events:
  - netconf_auth
  - netconf_unauth

  One of these two log events is sent upon receiving any RPC from a NETCONF client. The netconf_auth and netconf_unauth log events are mapped using the sros-md-rpc-accounting-event NETCONF notification.

  The NETCONF local command accounting log events and NETCONF notification do not show the details of the configuration changes in an <edit-config> RPC.

4.6.6 Model-driven Interfaces

- In Release 19.10.R1, configuration groups support for all containers in the **configure service vprn** *service-name* **interface** *interface-name* **sap** *sap-id* MD-CLI configuration branch is now ready for production networks. This feature was introduced in Release 16.0.R4.
- Release 19.10.R1 enhances configuration groups support for the following configuration branches and their descendants. This includes configuration groups definition and applying the groups with the **apply-groups** command.
  - Full configuration groups support: **aaa**, **cflowd**, **filter**, **lag** *lag-index*, **python**, **qos**, **subscriber-mgmt**
  - Partial configuration groups support (see the *MD-CLI User Guide* for a detailed list of exceptions): **card** *slot-number*, **port** *port-id*, **router** *router-name*, **service**, **system**

4.6.7 Configuration Coverage in Model-driven Interfaces

This section lists configuration commands that were unavailable in MD interfaces in previous releases, but are newly available for configuration in Release 19.10.R1. These are features and configuration commands that existed in classic CLI and SNMP in previous releases.

Refer to Unsupported Configuration in MD Interfaces for more details about configuration elements that are unsupported.

4.6.7.1 System

- Lawful Intercept (LI)
- **configure port transceiver** (CFP2-DCO optics)
- **configure mirror mirror-source**
- Ethernet satellites (including **system software-repository** and **system port-topology**)

4.6.7.2 QoS

- Hardware

  - IOM4-e-HS and HS-MDAv2 support (including all hsmda and hs qos policies and parameters)
- Card
  - FP egress **[NEW]**
- Port
  - Ethernet network egress queue group queue override
- QoS Policies
  - SAP egress IP/IPv6 entry action port redirection
  - Ingress queue group templates
  - IPv6 prefix lists
  - Post-policer mapping
  - Queue group redirect lists
- Router
  - Network domains
- Services
  - Applied SAP ingress QoS with the applied QoS policy configured with IPv6 or MAC criteria statements when used with subscriber management under VPLS SAPs, IES interface SAPs, VPRN interface SAPs
  - Scheduler policy applied to VPRN interface egress
  - **match-qinq-dot1q** under VPRN interface ingress SAPs
  - **qinq-mark-top-only** under
    - Epipe egress SAPs
    - VPRN interface egress SAPs

4.6.7.3 Multicast

- MSDP (router and VPRN)
- **mtrace2** configuration (router and VPRN)
- GTM

4.6.7.4 Layer 2 Services

- M-VPLS services (including **managed-vlan-list** and **mst-instance**)
- VPLS **mac-protect**

- ISID policy
- Static ISIDs
- B-VPLS

4.6.7.5 Subscriber Management

- **configure redundancy multi-chassis peer mc-ring l3-ring**

4.6.8 LAG

- In Release 19.10.R1, setting a specific IGP interface cost (under the **config**>**router**>**isis**>**if** or **config**>**router**>**ospf**>**if** contexts) no longer overrules the configured or calculated LAG cost values (**config**>**lag**>**port-threshold** or **weight-threshold** or **hash-weight-threshold**). In addition, Release 19.10.R1 no longer requires **reference-bandwidth** to be configured when only **static-cost** is used. [311576]
- Release 19.10.R1 introduces the capability to display LACP statistics per LAG port for packets transmitted and received using the **show lag** *lag-id* **lacp-statistics** command. [320419]

4.6.9 Ingress Multicast Path Management

- Release 19.10.R1 adds the **show**>**mcast-management**>**multicast-info-policy** command to display the **multicast-info-policy** information.[293364]

4.6.10 Routing

- In Release 19.10.R1, non-segmented MLDP Inter-as option B is extended to intra-AS (inter-area) option B. The intra-AS option B includes the support for ABR MoFRR.
- Release 19.10.R1 adds the following enhancements to the stitching of an SR-OSPF or SR-ISIS tunnel to an LDP FEC:

  - When IS-IS or OSPF resolves a node SID, it stitches the Segment Routing (SR) ILM to an LDP FEC when either there is no SID received with the prefix, or the SID which came with the prefix resolves to a non-SR enabled next-hop. The latter can be a next-hop over a network interface when the peer did not advertise the SR capability or over an IES or a VPRN interface.

    Prior to Release 19.10.R1, no attempt was made to stitch the SR tunnel to the LDP FEC of the same prefix of the SID received with the prefix, regardless of how it was resolved.
  - The stitching of the SR ILM to the LDP FEC does not require the presence of a Segment Routing Mapping Server (SRMS) SID entry for the prefix.

    Prior to Release 19.10.R1, implementations would stitch only if a SRMS SID entry for the prefix was received.

    When upgrading a configuration from a prior release, this behavior is enabled by default, and additional stitching of the SR ILM to the LDP FEC will be affected under the above conditions. [309355, 322023]
- In Release 19.10.R1, when an IS-IS or OSPF prefix is rejected due to the application of an import policy, the SR node-SID tunnel of the same prefix is also not programmed in the tunnel table and in data path.

  The behavior is modified so the import policy does not prevent the SR-ISIS or SR-OSPF tunnel to stitch to an LDP FEC of the same prefix when the next-hop is not SR-enabled. The only condition is there is an export policy which redistributes the prefix into the local IGP instance, which owns the SR-ISIS or SR-OSPF tunnel.

  The outcome of the import policy continues to apply when the SR-ISIS or SR-OSPF tunnel is normally resolved to a SR-enabled next-hop in IS-IS or OSPF.

  When upgrading a configuration from a prior release, this new behavior is enabled by default, and additional stitching of the SR ILM to LDP FEC is affected under the above conditions. [311691, 316266]
- Release 19.10.R1 relaxes the checks upon a configured IPv4 loopback address. The router now allows for a loopback IPv4 host address to overlap with the range configured on another configured IP router interface. Overlapping IPv6 addresses remain unsupported. [323692]
- Release 19.10.R1 introduces IPv6 forwarding adjacency support for RSVP-TE tunnels in single topology IS-IS. [325576]
- Release 19.10.R1 introduces two commands to allow control of the forwarding of IP-in-IP encapsulated traffic. The two new system level commands **config**>**system**>**ip**>**forward-6in4** and **config**>**system**>**ip**>**forward-ip-over-gre** enable the forwarding of IPv6-over-IPv4 and IP-over-GRE traffic and the default for both is disabled. In order to continue to forward these types of traffic, one or both of the new system level commands must be enabled.

In previous releases, IPv6-in-IPv4 and IP-over-GRE traffic directed to the SR OS system-IP address would, by default, be de-encapsulated and the inner IP packet would be forwarded to the appropriate destination.

These two new commands change the default behavior, but enhance security and provide greater control on the handling of encapsulated traffic. [326385]

- In Release 19.10.R1, when a node is in IS-IS overload state, it now keeps preferring local RSVP tunnel(s) instead of falling back to IGP route(s). [329916]
- Release 19.10.R1 extends the LFA policy feature to IPv4 SR-OSPF adjacency SID tunnels.

4.6.11 BGP

- Release 19.10.R1 improves the handling of BGP communities by removing duplicate values in the same attribute. This was already supported in some scenarios, but the scope is extended in Release 19.10.R1. [314574]
- Release 19.10.R1 extends QoS Policy Propagation via BGP (QPPB) and policy-accounting support for BGP label-IPv4 and label-IPv6 routes. [321885]

4.6.12 MPLS

- Local CSPF Path Computation for SR-TE LSP is now ready for production networks. This feature was introduced in Release 19.7.R1.
- S-BFD for SR-TE LSPs is now ready for production networks. This feature was introduced in Release 19.7.R1.
- Release 19.10.R1 introduces the ability to enable the configuration of the maximum number of Points of Local Repair (PLRs) per RSVP-TE bypass LSP. A PLR summarizes the constraints applied to the computation of the path of the bypass LSP. It consists of the avoid link/node constraint and potentially other TE constraints, such as exclude SRLG, that are needed to protect against the failure of the primary path of the RSVPTE LSP that is associated with the bypass LSP.

  Additional PLRs with the same avoid link/node constraint are associated with the same bypass to minimize the number of bypass LSPs created. This command controls the maximum number of such PLRs.

Since MPLS saves only the PLR constraints of the first LSP that triggered the dynamic bypass creation, subsequent LSPs for the same avoid link/node and with the non-strict bypass SRLG disjointness enabled may be associated with the same bypass. This happens even in cases where there exists a bypass LSP path that strictly satisfies the SRLG constraint. When the maximum PLRs per bypass is configured with a value of 1, MPLS triggers the signaling of a new dynamic bypass LSP for each new PLR and saves each PLR constraint separately with its own bypass. As a result, when MPLS re-optimizes a bypass LSP it guarantees that SRLG disjointness of that PLR are checked and enforced. [321633]

- Release 19.10.R1 increases the maximum bandwidth supported for RSVP-TE and SR-TE LSPs to 6.4Tbps. Bandwidth values of up to 6.4Tbps can be set for configured and auto LSP paths, as well as auto-bandwidth. The **tools**>**perform**>**router**>**mpls**>**cspf** and **tools**>**perform**>**router**>**mpls**>**adjust-autobandwidth** commands also support the increased maximum bandwidth. The increased maximum bandwidth is applicable to RSVP-TE and PCEP signaled LSPs. [322272]
- Release 19.10.R1 introduces the ad-hoc resignaling of all SR-TE LSPs at the receipt of one or more IGP link down events in TE-DB. Once the re-optimization is triggered, the behavior is exactly the same as the timer-based resignal or the delay option of the manual resignal. MPLS will force the expiry of the resignal timer and will ask TE-DB to re-evaluate the active paths of all SR-TE LSPs. The re-evaluation updates the total IGP or TE metric of the current path, checking the validity of the hops and labels, and computing a new CSPF for each SR-TE LSP. MPLS programs the new path, only if the total metric of the new computed path is different from the updated metric of the current path, or if one or more hops and labels of the current path are invalid. Otherwise, the current path is considered to be one of the most optimal ECMP paths and is not updated in data-path. [328682]
- Intra-AS Option B Support for MVPN is now ready for production networks. This feature was introduced in Release 19.5.R1.

4.6.13 Services

- Release 19.10.R1 extends the applicability of Class-based Forwarding with ECMP to VPRN-v4/v6 service prefixes which resolve to RSVP-TE LSPs.
- Release 19.10.R1 adds the support of XML accounting custom records for policers when the accounting policy record is configured to **custom-record-service**. The functionality is equivalent to that currently provided for custom records with queues with the exception that the counters collected are the intersection of the configured custom record counters and the statistics reported by the policer’s **stat-mode**.

- Release 19.10.R1 enhances the **tools**>**dump**>**resource-usage** command to include the SAP entry breakdown per service at the system level and the SAP instance breakdown per service at the line-card level.
- Release 19.10.R1 improves the support for PIM snooping in R-VPLS services where EVPN is enabled. In particular:
  - PIM snooping for IPv4 is now supported in R-VPLS services with EVPN-MPLS or EVPN-VXLAN destinations. The R-VPLS can be attached to a VPRN or IES service.
  - PIM snooping for IPv6 is now supported in R-VPLS services with EVPN-VXLAN destinations without MLD snooping. The R-VPLS can be attached to a VPRN or IES service.
- In Release 19.10.R1, Proxy-ARP/ND entries are now ESI-aware. The system responds to ARP/ND requests for a given IP1, when the ESI of the EVPN MAC/IP1 route and ESI for the MAC FDB Entry match, irrespective of the FDB Entry type (for example, static or dynamic). [330105]

4.6.14 Subscriber Management

- Release 19.10.R1 introduces the ability to configure a minimum reserved bandwidth as part of an egress aggregate rate limit under the subscriber profile. This is applicable to operators that use IGMP adjustment rates to subtract the total egress rate for the subscriber, where per-multicast channel bandwidth consumption is configured under the **config**>**router**>**mcac**>**bundle**>**policy** context. A household can consume all egress bandwidth with video traffic. This allows the BNG to ensure a minimum unicast bandwidth for the subscriber, regardless of the amount of bandwidth consumed by multicast. This function reserves a minimum bandwidth for emergency services, such as VoIP. Misuse of this feature can cause over-subscription to access bandwidth, affecting the end-user experience. This feature does not limit multicast bandwidth. Therefore, proper QoS provisioning is required on the access node to prioritize unicast traffic over multicast traffic. [327700]
- Release 19.10.R1 increases the maximum length of the subscriber and SLA profile strings from 16 characters to 32 characters. The subscriber and SLA profile strings are returned by dynamic interfaces (RADIUS and Diameter) and are used to select locally-configured subscriber and SLA profiles assigned to a subscriber session. Mapping between the subscriber and SLA profile strings and locally-configured subscriber and SLA profiles, which already support a 32-character length, can be direct or indirect through a mapping table in the subscriber-identification policy.
- In Release 19.10.R1, ESM-over-GTP no longer requires an additional internal host for each GTP subscriber.

- Release 19.10.R1 adds the support of XML accounting custom records for policers when the accounting policy **record** is configured to **custom-record-subscriber**. The functionality is equivalent to that currently provided for custom records with queues with the exception that the counters collected are the intersection of the configured custom record counters and the statistics reported by the policer’s **stat-mode**.

4.6.15 IPsec

- Release 19.10.R1 introduces the tIPsecTunnelProtocolFailed IPsec tunnel event which is generated when an abnormal IKE event (such as DPD timeout) occurs in an existing IPsec tunnel. [262819]
- Release 19.10.R1 improves the throughput of small number of IPsec tunnels with single CHILD_SA per tunnel on VSR.

4.6.16 NAT

- Release 19.10.R1 introduces the support for bridged homes with an **nh-mac** anti-spoof setting in ESM for L2aware NAT. This functionality is enabled using the **config**>**subscr-mgmt**>**sub-profile**>**nat-access-mode bridged** command.

  In this scenario, all bridged subscriber hosts are eligible for L2-Aware NAT, but with inferior anti-spoof in ESM. However, the source IP address of the upstream traffic is still checked by the NAT function and frames from the spoofed IP addresses are dropped. Spoofed IP addresses are considered IP addresses that do not belong to the bridged hosts initially set up in the ESM.

  This model has the following requirements model:

  - MAC addresses within the subscriber and SAP must be unique.
  - If this model is also deployed for routed RGs, the routed RGs must have NAT enabled. With this, the hosts behind the routed RG are hidden behind the RGs NAT and not visible in BNG.

4.6.17 Application Assurance

- Release 19.10.R1 introduces additional AA Cflowd configurable fields: [329612]
  - IP TTL (including reverse for comprehensive) volume and comprehensive templates:

- TCP ReTx bytes (including reverse for comprehensive)
- TCP ReTx Packets (including reverse for comprehensive)
- TCP SessEstDelay (The AA-measured TCP session delay between SYN and SYN_ACK)
- TCP SessEstDelay Reverse (The AA-measured TCP session delay between SYN_ACK and ACK)

4.6.18 Scaling

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

- The number of sessions for the **static-route-entry cpe-check**.
- The maximum number of egress port queue group instances per system, from 40k to 160k. This is not applicable to the HS-MDAv2 and IOM4-e-HS.
- The maximum number of egress network port queue group instances per card/slot, from 8k to 16k on FP2-, FP3- and FP4-based line cards, excluding the HS-MDAv2 and IOM4-e-HS.
- The maximum number of ingress policers, ingress policer statistics, ingress root arbiters, and egress policer statistics per-FP on FP4 hardware when using system profile none and system profile A.
- The MTU range has been increased to support up to 9800 octet frames on FP4 Ethernet ports. Routing and services MTUs have also been increased to support these larger frames. ISA applications and satellites MTUs have not been increased.
- The number of BGP optimal route reflector locations (reference points for SPF calculations).
- The maximum number of SAP instances per FP, per MDA and per Card on FP4 hardware in system profile none and A.
- The maximum number of SAP entries per system in system profile A.
- On CPM5, PIM SSM (S,G) can be scaled to 256K per system (in GRT) with a current limitation of per FP complex (S,G) scale. This scaling is done using the **config>router>pim>pim-ssm-scaling** command. The total maximum multicast capacity is constrained by the lowest performance FP.

  A PIM **shutdown** and **no shutdown** is needed to enable the scaling of (S,G)s.

4.7 Release 19.7.R2

4.7.1 Services

- EVPN-VXLAN Multi-homing Along with Non-system IPv4/IPv6 Termination is now ready for production networks. This feature was introduced in Release 19.7.R1.
- EVPN-VXLAN Multi-homing Support on R-VPLS is now ready for production networks. This feature was introduced in Release 19.7.R1.

4.8 Release 19.7.R1

4.8.1 Hardware

- Release 19.7.R1 introduces a **detail** option in the **show port description** command. This option adds the admin, link, and port state for each port. [275259]
- Release 19.7.R1 introduces a new **event** context under **card**>**mda**. An operator needs to first create the **event** and specify the event type as **soft-error**. The new **event** provides options on how to handle soft-errors on FP2- and FP3-based MDAs. The **event** can have the following actions to be taken on the MDA when a soft-error occurs: The actions are:
  - **no action** – Do nothing (basic error counting)
  - **log-only** – Generates a log-event, SNMP trap
  - **reset-mda** – Resets to clear the error
  - **fail-mda** – Fails the MDA until operator intervention

  The default action will be **log-only**. If the **event** is not created, no action will be taken. [313934]

4.8.2 System

- Release 19.7.R1 has been enhanced to perform faster polling of interface, port, and LAG information. This impacts MIB walks of the ifTable and ifXTable MIB tables. [319176]
- Release 19.7.R1 adds additional characters that can be used in a Syslog's log-prefix. Permissible characters are VCHAR values (%d33-126) except: "=", ":" and """. The additions are compliant with RFC 3164. [321852]
- Release 19.7.R1 is enhanced to support gnmi.proto version 0.7.0. All enhancements related to this upgrade are backwards compatible.
- IEEE 1588 Transparent Clock functionality on Ethernet satellites is now ready for production networks. This feature was introduced in Release 16.0.R5.

4.8.3 MD-CLI

- In Release 19.7.R1, dynamic defaults for parameters that change in relation to other configuration parameters are now displayed in the MD-CLI **?** help in addition to static defaults. For example, **configure log accounting-policy 1 collection-interval** has a dynamic default of 5 that changes to a dynamic default of 10 if **configure log accounting-policy 1 record video** is configured, and changes to a dynamic default of 15 if other **record** types are configured. [315240]
- Release 19.7.R1 improves the MD-CLI **show** command **?** help and command completion to be more flexible so that named parameter names are no longer required. [317089]

  For example, in previous releases the following commands required named parameters:

  - **show router bgp neighbor 1.1.1.1 filter1 received-routes**
  - **show router bgp neighbor 1.1.1.1 filter2 history**

  In Release 19.7.R1, the commands are:

  - **show router bgp neighbor 192.168.0.1 received-routes**
  - **show router bgp neighbor 192.168.0.1 history**

4.8.4 Model-driven Interfaces

- In Release 19.7.R1, the following IS-IS STATE leafs have been enhanced to support ON_CHANGE subscription and correspond to the following SNMP traps.

SNMP Trap		NOKIA YANG LEAF
2029 tmnxIsisDatabaseOverload	P0	Level overload status
2031 tmnxIsisCorruptedLSPDetected	P0	statistics corrupted-lsps
2035 tmnxIsisOwnLSPPurge	P0	Statistics own-lsp-purges
2038 tmnxIsisAuthFail	P0	statistics authentication-failures
2045 tmnxIsisAdjacencyChange	P0	Interface adjacency oper-state
2047 tmnxIsisAdjRestartStatusChange	P0	Interface adjacency restart status
2060 tmnxIsisLSPPurge	P0	Statistics own-lsp-purges

4.8.5 Configuration Coverage in Model-driven Interfaces

This section lists configuration commands that were unavailable in MD interfaces in previous releases, but are newly available for configuration in Release 19.7.R1. These are features and configuration commands that existed in classic CLI and SNMP in previous releases.

Refer to Unsupported Configuration in MD Interfaces for more details about configuration elements that are unsupported.

4.8.5.1 System

- NAT configuration
- DWDM wavetracker configuration

4.8.5.2 Routing

- Router **origin-validation**
- IP and GRE tunnels using ISA (including tunnel SAPs and interfaces)
- VRRP policies

4.8.5.3 Layer 2 Services

- **mac-move**
- VPLS **provider-tunnel**
- STP (under VPLS, SAP, spoke-SDP, PW-template)

4.8.5.4 QoS

- Services
  - Multi-service sites with a **policer-control-policy** applied to the ingress/ egress of a LAG SAP with the LAG in distribute mode and members spanning multiple FPs on an FP4 XMA

4.8.5.5 IPsec

- IPsec (including SeGW, **file-transmission-profile**, **isa tunnel-group**)

4.8.5.6 Subscriber Management

- Category Maps
- Layer 2 Access Points for WLAN-GW
- Python support for NAT and WLAN-GW groups
- Call trace
- L2-Aware NAT
- vRGW
- WLAN-GW distributed subscriber management
- IES and VPRN redundant interfaces
- **service template epipe-sap-template**

4.8.6 Ingress Multicast Path Management

- By default, the **configure mcast-management chassis-level per-mcast-plane-capacity total-capacity** command is set to **dynamic** to determine the capacity of IMPM paths and planes. In Release 19.7.R1, when **dynamic** is used with FP4 hardware in a 7750 SR-7-B/12-B chassis, the resulting **total-capacity** is reduced from 19000 to 17000.

4.8.7 Routing

- Release 19.7.R1 enhances the Cflowd MPLS-IP template behavior so that the MPLS label stack is always reported in the Cflowd flow data at the ingress and egress LER. This is supported for IES and VPRN services transporting either IPv4 or IPv6 traffic. In addition, sampling of IPv4 and IPv6 traffic that are being forwarded into an MPLS shortcut will also have their label stacks reported at both ingress and egress network sampling.

  This behavior change only affects flow data being sent to v9 or v10 Cflowd collectors. There are no changes to the MPLS-IP template, as the required fields already exist.

4.8.8 BGP

- Release 19.7.R1 adds a new optional **discard-component-communities** parameter to the **aggregate** command in both the Base and VPRN routing contexts. This new parameter causes the communities from contributing routes to not be aggregated into the newly-formed aggregate route. This command allows the explicitly configured communities from the **aggregate** route command itself to be advertised to outside BGP neighbors. [319771]
- Release 19.7.R1 extends the existing support of RPKI-based BGP prefix origin validation to VPRN BGP instances. When origin validation is enabled on a VPRN PE-CE BGP session, the IP routes received from the peer are assigned origin validation states based on the lookup of each route’s origin AS and IP prefix in the router’s database of origin validation entries. This database is shared by all VPRN BGP instances and the Base router BGP instance. It consists of entries learned from static configuration and RPKI and router protocol interaction with local cache servers.

- In Release 19.7.R1, BGP import policies and VRF import policies can change the next-hop address in any received BGP route to an IPv6 address. In previous releases, import policies could not change the next-hop to an IPv6 address in the routes of some address families.

  Also in Release 19.7.R1, the action in VRF export policies to set the BGP next-hop to an IPv6 address is now honored for VPN-IPv4 routes that are advertised to an RFC 5549-capable peer. [321887]
- After an upgrade to Release 19.7.R1, **as-override** will now replace a BGP peer's autonomous system number (ASN) in the AS-path with the **local-as** number instead of the configured **autonomous-system** number if a different **local-as** number is configured on this peer while advertising routes to it. [323787]

4.8.9 PIM

- Prior to Release 15.0, the RPF check for NG-MVPN was done against the IPv4 VPN address-family (AF) next-hop field. After the introduction of inter-AS non-segmented MLDP in Release 15.0, the RPF check was done against the VRF import extended community field in the IPv4 VPN route. This new behavior is compliant with RFC 6514. If using a BGP router ID that is different from the system IP address, and deploying NG-MVPN, the VRF route import extended community will be set to this non-system IP address. In addition, by default, the MVPN advertises its AD routes with the system IP address; therefore, interoperating between a Release 15.0 SR OS node and an SR OS node running Release 14.0 or lower will fail as the RPF check will fail. Releases 15.0 and higher used the VRF import extended community.

  To allow an upgrade path, Release 19.7.R1 introduces a new **apply-bgp-nh-override** CLI option in the **configure service vprn pim** context to force the RPF to be checked against the IPv4 VPN AF next-hop and not the VRF route import extended community. By default, this option is disabled and the RPF is checked against VRF route import. In Release 19.7.R1, this option is only available in MVPN. [318340]

4.8.10 Subscriber Management

- Release 19.7.R1 enhances the Home LAN Extension (HLE) with a traffic policer function enabled using the **config**>**service**>**ies|vprn**>**sub-if**>**grp-if**>**wlan-gw**>**ranges**>**range**>**vrgw**>**lanext**>**access|network**>**policer** command. An ISA policer can be specified for ingress tunnel traffic forwarded on following HLE connections:

  - Access (Home)
  - Network (DC)
- In Release 19.7.R1, the concatenated length for each of the following RADIUS attributes has been increased:
  - from 494 to 1235 bytes: Alc-ToServer-Dhcp-Options and Alc-ToServer-Dhcp6-Options
  - from 494 to 1729 bytes: Alc-ToClient-Dhcp-Options and Alc-ToClient-Dhcp6-Options

  This feature is supported for both DHCP relay and proxy. Fragmented DHCP packets are not supported. The port MTU should be increased accordingly when DHCP packets larger than 1500 bytes are expected.

4.8.11 IPsec

- Release 19.7.R1 enhances IPsec reverse route creation for dynamic LAN-to-LAN tunnel to allow the system to ignore full address ranges (for example, 0.0.0.0 to 255.255.255.255) in Traffic Selector initiator (TSi) when creating a reverse route. This behavior is optional and can be enabled using the **config**>**ipsec**>**tnl-temp**>**sp-reverse-route ignore-default-route** command.

4.8.12 NAT

- In Release 19.7.R1, for MD Interfaces, **map** commands in deterministic NAT and 1:1 agnostic NAT must be fully specified by the operator (or client). To simplify this process in the MD-CLI for cases where custom mappings are not required, a new **tools perform nat deterministic calculate-maps** command is introduced. This command outputs a set of system-generated **map** statements in an MD-CLI configuration format. The **map** commands can then be copied and pasted into an MD-CLI candidate by the operator. This new **tools** command offers some configuration assistance for operators using manual configuration of NAT in the MD-CLI.

4.8.13 OAM

- Release 19.7.R1 enhances the **show oam-pm sessions** command on the OAM-PM **mpls-dm** querier to display a list of **mpls-dm** error conditions reported by the responder. The augmentation **show oam-pm sessions test-family mpls detectable-rx-errors** command is specific to the MPLS **test-family**. [318483]

4.8.14 Scaling

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

- On CPM5s, PIM SSM (S,G) per system (in GRT) with current limitation of per-FP complex (S,G) scale. This scaling is possible using the **config**>**router**>**pim**>**pim-ssm-scaling** command. When **pim-ssm-scaling** is enabled, ASM, DM, MoFRR, JP Policy and SSM groups are not supported. This feature works with IMPM.
- The maximum number of IPsec tunnels per VSR instance

4.9 Release 19.5.R2

4.9.1 Services

- IPv4 MPLS Forwarding Policy and IPv4 RIB API Tunnel Resolution for EVPN MPLS Services are now ready for production networks. This feature was introduced in Release 19.5.R1.

4.10 Release 19.5.R1

4.10.1 Hardware

- Release 19.5.R1 adds the support for tmnxEqCardPChipError events on HS-MDAv2. [252341]
- In Release 19.5.R1, the 7750 SR-7s/14s active CPM can communicate with the power shelves through the RS-485 cable connected to the PWR port on either CPM. The standby CPM must be online and in sync with the active CPM. When an RS-485 cable is disconnected or missing from either of the CPMs, or if there are two power shelves and the interconnecting RS-485 cable is missing, the active CPM will proxy-communicate through the standby CPM. A minor alarm is generated to alert of a situation where a power shelf is not reachable through one or both CPMs. [310238]
- IEEE 1588 Port-Based Timestamping (PBT) on 7750 SR-1s systems is now ready for production networks. This feature was introduced in Release 16.0.R4.

4.10.2 System

- In Release 19.5.R1 the default of the SR OS **login-banner** has been changed to avoid signaling any information related to product or software release prior to a successful authentication. This enhancement is most visible via SSH/Telnet to use CLI/NETCONF to manage an SR OS router. The option to display a login-banner can be controlled using the **configure system login-control login-banner** command. [314675]

4.10.3 Model-driven Interfaces

- In Release 19.5.R1, the **eth-cfm** container in the **configure service vprn** *service-name* **interface** *interface-name* **sap** *sap-id* branch of MD-CLI configuration groups is ready for production networks. This feature was introduced in Release 16.0.R4.
- Release 19.5.R1 adds the configuration groups support in model-driven mode for the following configuration branches and its descendants. This includes configuration groups definition and applying the groups with the **apply-groups** command.

  - **configure card** *slot-number* **fp** *fp-number*
  - **configure service vpls** *service-name*
- Release 19.5.R1 includes configuration group definitions and **apply-groups** commands in the YANG model, enabling the use of configuration groups via NETCONF and gNMI. The NETCONF <get-config> operation and the gNMI Get RPC command returns the pre-expanded configuration, including the configuration groups definitions and **apply-groups** configuration elements. The expanded configuration, including inherited configuration elements, cannot be returned with NETCONF or gNMI.

4.10.4 Configuration Coverage in Model-driven Interfaces

This section lists configuration commands that were unavailable in MD interfaces in previous releases, but are newly available for configuration in Release 19.5.R1. These are features and configuration commands that existed in classic CLI and SNMP in previous releases.

Refer to Unsupported Configuration in MD Interfaces for more details about configuration elements that are unsupported.

Release 19.5.R1 supports:

- PTP (IEEE 1588) in mixed configuration-mode only, using classic interfaces only
- Additional **sonet-sdh** configuration commands for use in 10G Ethernet WAN mode (xgig): section-trace and threshold
- port ethernet crc-monitor, dampening, and symbol-monitor
- DNSSEC
- Additional L2 services coverage:
  - **proxy-arp**
  - **proxy-nd**
  - BPDU translation
  - L2PT
  - **bgp-vpls**
  - **static-mac** (except for SPB-enabled services)
- Additional L3 services coverage:
  - **configure service vprn vxlan tunnel-termination**
- Multicast: MVPN, BIER
- TPSDA: Enhanced and Basic Subscriber Management (BNG and WLAN-GW)

- non-ESM/BSM DHCP relay
- L2oGREv4/v6 (FPE based, including **gre-eth-bridged** SDPs)
- BGP **dynamic-neighbor**
- Additional QoS coverage:
  - SAP bandwidth CAC
  - SAP ingress QoS policy FC assignment, except for FP ingress queue group redirection
  - Mirror destination SAP egress QoS
  - Egress port scheduler overrides
  - Epipe SAP egress scheduler policy
  - VPLS SAP ingress and egress:
    - policer overrides
    - policer control policy overrides
  - IES interface SAP egress:
    - queue overrides
    - policer overrides
    - scheduler policy
    - scheduler overrides
    - policer control policy overrides
  - IES interface SAP ingress:
    - queue overrides
    - policer overrides
    - scheduler overrides
    - policer control policy overrides
  - VPRN interface SAP egress:
    - policer overrides
    - policer control policy overrides

4.10.5 NETCONF

- Release 19.5.R1 adds the support for the following <copy-config> scenario: <source>=<config> and <target>=<candidate>. This is a combination that can be used by operators to copy golden configurations to a new or existing SR OS box.

4.10.6 Telemetry

- Release 19.5.R1 allows the minimum sampling interval of one (1) second to be specified. The default sampling interval remains ten (10) seconds.

4.10.7 Routing

- Release 19.5.R1 increases the size of the **config**>**router**>**if**>**description** *long-description-string* field from 160 to 255 characters.
- BIER is now ready for production networks. This feature was introduced in Release 16.0.R4.
- TCP Authentication Option (TCP-AO) support for LDP is now ready for production networks. This feature was introduced in Release 16.0.R4.

4.10.8 BGP

- Release 19.5.R1 increases the maximum number of BGP standard communities, extended communities, and large communities that can be present in one BGP Update message. Prior to Release 19.5.R1, the related path attributes were considered too long if any one of them exceeded 2048 bytes in length. This caused associated routes to be marked as invalid. [303820]
- In Release 19.5.R1, the EBGP and IBGP default route preference can be configured as separate values. In previous releases, it was necessary to use route policies to make the preference different in all EBGP routes versus all IBGP routes. [307883]
- In Release 19.5.R1, locally defined aggregate routes are no longer automatically preferred over other BGP routes for the same prefix. They now only win over another BGP route if they have the same or better route preference (numerically equal or lower). With default route preferences for BGP and aggregate routes, no change of behavior will be observed. [313975]

4.10.9 BGP-EVPN

- In Release 19.5.R1, the AC-DF bit in the DF Election extended community advertised along with the EVPN ES routes is correctly set. This correction is per RFC 8584. [308142]

4.10.10 MPLS

- Release 19.5.R1 extends the PCEP PCE implementation by adding the support of a path update with PCC implementations that do not follow the Make-Before-Break (MBB) procedure. When following the MBB procedure, the PCC must change the LSP-ID in the LSP-Identifiers TLV when sending a report message (PCRpt) to PCE following a path update message (PCUpd).

  RFC 8231 does not explicitly exclude the use of a non-MBB procedure by which the PCC reports the same LSP-ID for the new path in the PCRpt message. The VSR-NRC PCE now supports the non-MBB procedure to interoperate with third-party implementations of the PCEP PCC function. [312385]

4.10.11 LDP

- Release 19.5.R1 extends routing policy support to be assigned to MLDP as an import policy. When a routing policy is assigned to MLDP as an import policy, the arriving label mapping (FEC) remains unresolved, or is resolved based on the policy prefix and action.

4.10.12 QoS

- In Release 19.5.R1, the configuration of buffer pool information and network queues has been moved from the **card mda network ingress** CLI tree to the **card fp ingress network** CLI tree. This enhancement enables per-FP configuration which is applicable to multi-FP XMAs. An **fp** parameter has been added to the **show pools** command, and the output has been re-organized to improve readability.
- Release 19.5.R1 increases the maximum configurable rate to 6400 Gb/s for the following items:
  - SAP ingress and egress QoS Policy queue and policer PIR and CIR rates
  - Scheduler policy scheduler PIR and CIR rates
  - Port scheduler max rate, group PIR and CIR rates and level PIR and CIR rates
  - Policer control policy root max rate and arbiter rate
  - SAP egress, multi-service site and Vport aggregate rates
  - Port and SAP overrides of the above except for:
    - Policer control policy root max rate overrides for queue groups

- Policer PIR and CIR rate overrides for queue groups
- Queue PIR and CIR rate overrides for queue groups

If the rates at ingress exceed the port capacity, or exceed the FP capacity with **per-fp-ing-queuing** configured, the rates are set to **max**. At egress, if the rates exceed the port capacity (including the **egress-rate** setting), they are set to **max**.

Rates greater than the above (capped) rates are only relevant when configured on a distributed or port-fair-mode LAG spanning multiple FPs.

The maximum bandwidth used depends on the card type. Consequently, the maximum rate used can change and the behavior of some existing configurations would also change. This also impacts the use of the **percent-rate** command with no parent or a *max-rate* parent, or the use of the **advanced-config-policy** where a **percent** *percent-of-admin-pir* is used.

Note that due to the changes in this implementation, there may be small differences in the resulting rates and thresholds compared to the previous implementation.

4.10.13 Services General

- Release 19.5.R1 enhances the MCS protocol to allow the optional enabling of transport encryption of synchronized IPsec states, which further increases the security of an MC-IPsec deployment. The **config**>**system**>**keychain** command is used to specify the key for encryption.

4.10.14 Application Assurance

- In Release 19.5.R1, HTTP error-redirect is enhanced to support redirect on the 400-417, 421-431, 451, 500-504, 505-511, 730, 731, and 735 error codes. [284732]
- Release 19.5.R1 adds the support to redirect HTTPS pages to a landing portal. Similar to HTTP Redirect, an HTTP Redirect object and an AQP can be configured to perform HTTPS redirection. If traffic matches to the configured AQP, AA returns a Nokia-signed certificate and traffic is redirected to the informative page. [307435]
- In Release 19.5.R1, AA Cflowd is enhanced to allow operators to select which fields to include in the exported Cflowd templates. [307539]

- In Release 19.5.R1, the AA Maximum Segment Size (MSS) adjustment feature is enhanced to support IPv6 packets marked as fragmented packets (atomic fragments), and fragmented packets where the MSS is in the first fragment. [316500]
- In Release 19.5.R1, RTP performance measurements adds the support for additional new VOIP codecs (such as AMR-WB+, EVS, SIREN14, IREN22, EVRCB, EVRNW, EVRCWB, etc).

4.10.15 Scaling

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

- The maximum number of 40GE, 100GE, and 400GE ports per LAG.
- The maximum number of IPoE sessions per BRG instance in vRGW.
- The BFD session scale on the 7750 SR-7s/14s.
- When dedicated to a single purpose TWAMP Light Reflector function, the 7750 SR-a4/a8 packet-per-second reflection of arriving TWAMP Test packets.
- The maximum bandwidth per LAG and the maximum number of 40GE, 100GE, and 400GE ports supported per LAG.
- The **lag**>**access**, **port**>**ethernet**>**access**, **service**>**ies**|**vprn**>**interface**>**sap**, and **service**>**vpls**|x**pipe**>**sap** bandwidth values. This increased value range allows meaningful access bandwidth values for higher-speed LAGs.

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 9 **Limited Support Features**

Section	Feature	Release Introduced
Hardware	virtualized Simulator (vSim)	12.0.R4
	Bluetooth Console Interface	16.0.R1
Services	Perfect Stream	8.0.R4
	Ad Insertion (ADI)	8.0.R4
	MACsec XPN modes	15.0.R5
	MLDv2-over-IPsec	15.0.R8

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 19.10.R6 Supported Hardware and Usage Notes for more information.

6.1 Hardware

Table 10 **Unsupported Hardware Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-1e/2e/3e	7450 ESS	7750 SR-1	7750 SR-1s/2s/7s/14s
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 [1]	X	X
Power-save Mode				X [2]	X [3]		X	X [4]

Notes:

1. VSM is only supported using the 7750 SR VSM-CCA-XP module in a 7750 SR IOM.
2. Not supported on 7750 SR-a4.
3. Not supported on 7750 SR-1e.
4. Not supported on 7750 SR-1s.
5. Supported on the 7750 SR-2s, but not the 7750 SR-1s/7s/14s.

6.2 System

Table 11 **Unsupported System Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-1e/2e/3e	7450 ESS	7750 SR-1	7750 SR-1s/2s/7s/14s
IEEE 1588 PTP	[1]							
IEEE 1588 Port-Based Timestamping (PBT)	[2]							
ACL Filter Egress Rate-Limit Action				X				
BITS input port redundancy							X [3]	X [3]
Centralized (CPM-based) CPU-Protection				X [4]	X [4]		X [4]	X [5]
Forwarding Path Extension (FPE) IP GRE Tunnel without ISA Port Cross-Connect (PXC) [6]								
Ingress Multicast Path Management				X	X		X	
Major ISSU Across Two Major Releases [7]							X	X
Minor ISSU							X	[13]
OOB Management Ethernet Port Redundancy				X	X		X	X [13]
Ethernet Satellites [8]		[10]	[10]					
SONET/SDH Satellites [9]	X					X	X	X
Soft Reset							X	[13]
System Alarm Contact Inputs	X	X	X		X	X	X	X
400GE LAG	X	X	X	X	X	X	X	[12]
Zero Touch Provisioning (ZTP) for IPv4 and IPv6 Networks	X	X	X	X	[14]	X		

Notes:

1. Not supported on the 7950 XRS-16c; supported on the 7950 XRS-20/20e and XRS-40.
2. 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.
3. BITS Input is supported, but there is no redundancy on a 7750 SR-1 or 7750 SR-1s.
4. Note that Distributed CPU Protection (DCP) is supported.
5. Supported on the 7750 SR-7s/14s but not supported on the 7750 SR-1s/2s.
6. Refer to Usage Notes for information on PXC platform support.
7. Not supported on satellites.
8. 7210 Ethernet satellites use 7210 SAS Release 9.0, 10.0, or 11.0. The 10GE Ethernet satellite requires 7210 SAS Release 9.0.R4 or higher. The Mxp satellite requires 7210 SAS Release 10.0.R5 or higher.
9. 7210 SONET and SDH satellites use 7705 SAR Release 8.0.R4 or higher.
10. Not supported on IOM4-e-HS.
11. Not supported on FP4-based hardware assemblies.
12. Requires QSFP-DD-capable XMA/system.
13. Not supported on the 7750 SR-1s/2s, but supported on the 7750 SR-7s/14s.
14. Supported on the 7750 SR-1e, but not the 7750 SR-2e/3e.

6.3 Quality of Service

Table 12 **Unsupported QoS Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-1e/2e/3e	7450 ESS	7750 SR-1	7750 SR-1s/2s/7s/14s
Named Pools	X	1, 3	1, 3	X	X		X	X
Ingress shared queuing (Dual-Pass)	X	1, 3	1, 2, 3				X	X
Policers (except for Distributed CPU Protection)				X				
H-QoS managed policers	X							X

Notes:

1. Not supported on IOM4-e-HS.

2. Not supported on 400G FP3 line cards.
3. Not supported on FP4 line cards.

6.4 MPLS

Table 13 **Unsupported MPLS Features**

Feature	**7950 XRS**	**7750 SR-7/12**	**7750 SR-12e**	**7750 SR-a4/a8**	**7750 SR-1e/2e/3e**	**7450 ESS**	**7750 SR-1**	**7750 SR-1s/2s/7s/14s**
GMPLS UNI		[1]	[1]			X		

Note:

1. Not supported on IOM4-e-HS.

6.5 Services

Table 14 **Unsupported Services Features**

Feature	**7950 XRS**	**7750 SR-7/12**	**7750 SR-12e**	**7750 SR-a4/a8**	**7750 SR-1e/2e/3e**	**7450 ESS**	**7750 SR-1**	**7750 SR-1s/2s/7s/14s**
Circuit Emulation services (for example, Cpipe SAPs)	X			X	X		X	X
new-qinq-untagged-sap configurability for :*.0 and :0.0 SAPs	X [1]							
FCC/RET/VQM	X			X	X	X	X	X

Table 14 Unsupported Services Features (Continued)

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-1e/2e/3e	7450 ESS	7750 SR-1	7750 SR-1s/2s/7s/14s
Frame Relay interfaces and services (for example, Fpipe SAPs)	X			X	X		X	X
config mirror mirror-source				X	X			
Tunnel services (IPsec, GRE, IP-in-IP tunnel termination) [2, 3]	X			X			X	X
IPv6 tunnel services (IPsec, GRE, IP-in-IP tunnel termination) [2, 3]	X			X			X	X
Throttling IKE messages	X			X			X	X
G.8031 (Ethernet tunnel support)		6	6					
Multi-Chassis features using IPsec (MC-IPsec) [4]	X			X			X	X
sFlow		6	6	X	X	X		
ECMP for VXLAN IPv4 Tunnels of R-VPLS [5]								

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. Requires an MS-ISA/MS-ISA2/MS-ISM.
5. All of the cards in the system must be FP3-based.
6. Not supported on IOM4-e-HS.

6.6 Subscriber Management

Table 15 **Unsupported Subscriber Management Features**

Feature	**7950 XRS**	**7750 SR-7/12**	**7750 SR-12e**	**7750 SR-a4/a8**	**7750 SR-1e/2e/3e**	**7450 ESS**	**7750 SR-1/1s/2s/7s**	**7750 SR-14s**
IPv4 local DHCP Server	X							X
IPv6 local DHCP Server	X							X
GTP Uplink	X			X		X		X
L2TP LNS [1, 2]	X			X			X	X
port-policy command [1, 2]	X			X			X	X
NAT [1, 2]	X			X			X	X
Subscriber Accumulated Statistics	X					X		X
Subscriber Management—Routed CO (VPRN/IES subscriber interfaces)	X							X
Subscriber Management—Bridged CO (VPLS)	X							X
vRGW on regular group interfaces [1, 2]	X			X		X	X	X
vRGW on WLAN-GW group interfaces [1, 2]	X			X		X	X	X
WLAN gateway (WLAN-GW) [1, 2]	X			X		X	X	X
GTP Access [3]	X			X		X		X
Bonding [3]	X			X		X		X
Oversubscribed Multi-Chassis Redundancy (OMCR)	X			X	X	X	X	X
Call-trace	X						[4]	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. Requires FPE. See Table 11.
4. Not supported on 7750 SR-1 and SR-1s.

6.7 Application Assurance

Table 16 **Unsupported AA Features**

Feature	7950 XRS	7750 SR-7/12	7750 SR-12e	7750 SR-a4/a8	7750 SR-1e/2e/3e	7450 ESS	7750 SR-1	7750 SR-1s/2s/7s/14s
Application Assurance [1]	X			X			X	X
AARP	X			X			X	X

Note:

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.

6.8 System Profiles

In Releases 16.0.R1 and higher, SR OS supports multiple system profiles which provide flexibility by enabling 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.

System profile none allows FP3- and FP4-based line cards to co-exist within a system. System profiles **profile-a** and **profile-b** support only FP4-based line cards.

Scaling numbers may differ per system profile. Contact your Nokia representative for details.

This section provides details of the unsupported features, and related commands, for a specific system profile:

- When no system profile is configured, or system profile **profile-a** is configured, the following features are not supported:
  - Segment Routing traffic statistics
  - Accounting for dark bandwidth
- When system profile **profile-b** is configured, the following features are not supported:
  - Application Assurance
  - Cpipe services
  - Dynamic data services
  - Subscriber Management
  - Ipipe services
  - Ipipe spoke-SDP termination on IES and VPRN

7 Obsoleted Features and Hardware

The following sections describe features and hardware that are no longer supported in SR OS. Obsoleted features from Releases 16.0.R2 to 16.0.R7 also apply to Release 19.*x*. Refer to the most recent *SR OS 16.0 Release Notes* for the summary of obsoleted features in Releases 16.0.R2 through 16.0.R7.

Note:

- The release image should not be loaded onto platforms that are no longer supported.
- Obsoleted hardware should be removed from the router before upgrading.
- Obsoleted features should be deconfigured on the router before upgrading.

7.1 Release 19.10.R6

In Release 19.10.R6, no features or hardware are obsoleted.

7.2 Release 19.10.R5

In Release 19.10.R5, no features or hardware are obsoleted.

7.3 Release 19.10.R4

In Release 19.10.R4, no features or hardware are obsoleted.

7.4 Release 19.10.R3

In Release 19.10.R3, no features or hardware are obsoleted.

7.5 Release 19.10.R2

In Release 19.10.R2, no features or hardware are obsoleted.

7.6 Release 19.10.R1

In Release 19.10.R1, no features or hardware are obsoleted.

7.7 Release 19.7.R2

In Release 19.7.R2, no features or hardware are obsoleted.

7.8 Release 19.7.R1

In Release 19.7.R1, no features or hardware are obsoleted.

7.9 Release 19.5.R2

In Release 19.5.R2, no features or hardware are obsoleted.

7.10 Release 19.5.R1

7.10.1 Hardware

7.10.1.1 7750 SR-c4/c12

The 7750 SR-c4 and SR-c12 chassis are obsoleted in Release 19.5.R1.

7.10.1.2 Obsoleted SF/CPMs, CFMs, MDAs, CMAs, MCMs, and CCMs

The hardware assemblies listed in Table 17 are no longer supported in SR OS starting in Release 19.5.R1.

Table 17 **Obsoleted SF/CPMs, CFMs, MDAs, CMAs, MCMs, and CCMs**

Nokia Part #	Description	CLI Name
3HE00023AA	20-port 100FX MDA – SFP	m20-100eth-sfp
3HE00030AA	1-port 10GBASE-LW/LR MDA with optics – Simplex SC	m1-10gb
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
3HE00231AA	20-port 100FX MDA – SFP	m20-100eth-sfp
3HE00235AA	1-port 10GBASE-LW/LR MDA with optics – Simplex SC	m1-10gb
3HE00709AA	7750 SR 1-port OC-192c/STM-64c MDA with LR-2/L-64.2 optic – Simplex SC	m1-oc192
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
3HE02185AA	2-port OC-3c/STM-1c/OC-12c/STM-4c CMA – SFP	c2-oc12/3-sfp
3HE03077AA	1-port Channelized OC-3/STM-1 CES CMA	c1-choc3-ces-sfp
3HE03607AA	7750 SR-c12 CFM-XP	cfm-xp
3HE03608AA	7750 SR-c4/c12 MCM-XP	mcm-xp
3HE03609AA	1-port GE SFP – CMA-XP	c1-1gb-xp-sfp
3HE03610AA	5-port GE SFP – CMA-XP	c5-1gb-xp-sfp
3HE03617AA	7750 SR-12 SF/CPM3	sfm3-12
3HE03618AA	7450 ESS-12 SF/CPM3	sfm3-12
3HE04164AA	7750 SR-7 SF/CPM3	sfm3-7
3HE04166AA	7450 ESS-7 SF/CPM3	sfm3-7

Table 17 **Obsoleted SF/CPMs, CFMs, MDAs, CMAs, MCMs, and CCMs (Continued)**

Nokia Part #	Description	CLI Name
3HE04580AA	7750 SR-c12 CCM-XP	ccm-xp
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
3HE05948AA	7750 SR-12 SF/CPM4	sfm4-12
3HE05949AA	7750 SR-7 SF/CPM4	sfm4-7
3HE05950AA	7450 ESS-12 SF/CPM4	sfm4-12
3HE05951AA	7450 ESS-7 SF/CPM4	sfm4-7
3HE06521AA	2-port OC-3c/STM-1c/OC-12c/STM-4c CMA – SFP Rev B	c2-oc12/3-sfp-b
3HE00709AA	1-port OC-192c/STM-64c MDA with LR-2/L-64.2 optic – Simplex SC	m1-oc192
3HE07166AA	7750 SR-12e SF/CPM4-12e	sfm4-12e
3HE07167AA	7750 SR-12e Mini-SFM4-12e	m-sfm4-12e
3HE08173AA	7750 SR-c12 CFM-XP-B	cfm-xp-b
3HE08220AA	8-port Channelized DS1/E1 CMA Rev B	c8-chds1

7.10.2 Satellites

- In Release 19.5.R1, the support for 7210 Ethernet satellites running 7210 SR OS Release 8.0 has been removed. Ethernet satellites running 7210 SR OS Release 8.0 must be upgraded to a supported release (Release 9.0, 10.0, or 11.0) either before or as part of the 7450 ESS, 7750 SR, or 7950 XRS upgrade process to Release 19.5.R1.

  For details on upgrading the Ethernet satellite software version refer to the “Satellite Software Upgrade Overview” section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*.

7.10.3 Filters

- The capability to test the reachability of Redirect Policies destinations by means of SNMP or URL based tests is removed in SR OS Release 19.5.R1. The commands (and all sub commands) of **config**>**filter**>**redirect-policy**>**destination**>**snmp-test** and **config**>**filter**>**redirect-policy**>**destination**>**url-test** are removed.

  Redirect Policies propose enabling two other types of tests to achieve an equivalent functionality: the ping test and unicast route test. These tests do not provide the exact same capability as they do not allow for the destination priority to be changed based on test result. [305518]

7.10.4 Subscriber Management

- ESM PPPoA/PPPoEoA is no longer supported as of Release 19.5.R1.

7.10.5 VSR

- VSR-D is no longer supported as of Release 19.5.R1.

8 Changed or Removed Commands

This section describes the SR OS commands that have been changed or removed in classic CLI or SNMP and commands that have been changed or removed in model-driven interfaces. Unless otherwise noted, all changed and removed CLI commands are accepted during boot and are converted on upgrade.

In the classic CLI, removed or changed commands are accepted in a script called with the **exec** command. In the MD-CLI, removed or changed commands are accepted (and converted as needed) in a configuration loaded with the **load full-replace** command.

This section contains:

- Changed or Removed Commands in Classic Interfaces
- Changed or Removed Commands in Model-driven Interfaces

AAA security profile rules (command authorization) that match on removed elements are not converted during an upgrade. Operators must examine their AAA security profile rules (locally on the router, or remotely on RADIUS or TACACS+ servers) and adjust them for any relevant command changes.

See also Software Upgrade Procedures for more information about the behavior of commands or parameters that have been modified or removed between releases. Changed or removed commands from Releases 16.0.R2 to 16.0.R7 also apply to Release 19.*x*. Refer to the most recent *SR OS 16.0 Release Notes* for the summary of changed or removed commands in Releases 16.0.R2 through 16.0.R7.

8.1 Changed or Removed Commands in Classic Interfaces

8.1.1 Release 19.10.R6

In Release 19.10.R6, no commands are changed or removed in classic interfaces.

8.1.2 Release 19.10.R5

In Release 19.10.R5, no commands are changed or removed in classic interfaces.

8.1.3 Release 19.10.R4

In Release 19.10.R4, no commands are changed or removed in classic interfaces.

8.1.4 Release 19.10.R3

In Release 19.10.R3, no commands are changed or removed in classic interfaces.

8.1.5 Release 19.10.R2

In Release 19.10.R2, no commands are changed or removed in classic interfaces.

8.1.6 Release 19.10.R1

8.1.6.1 NETCONF

- In Release 19.10.R1, the default value of the **config**>**system**>**management-interface**>**yang-modules**>**base-r13-modules** command is changed to **config**>**system**>**management-interface**>**yang-modules**>**no base-r13-modules**.
- In Release 19.10.R1, the default value of the **config**>**system**>**netconf**>**capabilities**>**writable-running** command is changed to **config**>**system**>**netconf**>**capabilities**>**no writable-running**.

8.1.6.2 MPLS

- In Release 19.10.R1, the **config**>**router**>**mpls**>**sr-te-resignal-timer** *minutes* command (which was released in a limited support state in Release 19.7.R1) is replaced with the **config**>**router**>**mpls**>**sr-te-resignal**>**resignal-timer** *minutes* command under a new node. The older **sr-te-resignal-timer** *minutes* command must be removed before upgrading to Release 19.10.R1.

8.1.6.3 Subscriber Management

- Release 19.10.R1 increases the length of the subscriber and SLA profile strings from 16 characters to 32 characters. See Enhancements in Release 19.10.R1 for more information. The following entries are affected by this change:

  RADIUS attributes:

  ```
  Alc-SubscProf-Str  (attribute id 26.6527.12)
  Alc-SLAProf-Str  (attribute id 26.6527.13)
  ```

  Gx attributes:

  - Charging-Rule-Name (AVP ID 1005)
    - Sla-Profile: *sla-profile-string*
    - Sub-Profile:*sub-profile-string*

  The *sla-prof-string* and *sub-prof-string* in the following CLI commands:

  ```
  configure subscriber-mgmt sub-ident-policy <name> sla-profile-map
       use-direct-map-as-default
       entry key <sla-prof-string> {
            sla-profile <sla-profile-name>
       }
  configure subscriber-mgmt sub-ident-policy <name> sub-profile-map
       use-direct-map-as-default
       entry key <sub-prof-string> {
            sub-profile <sub-profile-name>
       }
  configure subscriber-mgmt sub-profile <name>
       sla-profile-map
            use-direct-map-as-default
            entry key <sub-prof-string> {
                 sub-profile <sub-profile-name>
            }
  configure subscriber-mgmt local-user-db <name>
       ipoe
            host <name>
                 identification-strings
                      sla-profile-string <sla-profile-string>
                      sub-profile-string <sub-profile-string>
       ppp
             host <name>
                 identification-strings
  ```

```
sla-profile-string <sla-profile-string>
sub-profile-string <sub-profile-string>
```

MIBs containing sub-profile and sla-profile strings:

- tmnxSubProfSLAProfileMapTable
    - tmnxSubProfSLAProfMapSLAString
- tmnxSubIPolSLAProfileMapEntry
    - tmnxSubIPolSLAProfMapSLAString
- tmnxSubIPolSubProfileMapEntry
    - tmnxSubIPolSubProfMapSubString
- TmnxLocUsrDbPppoeEntry
    - tmnxLocUsrDbPppoeSlaProfString
    - tmnxLocUsrDbPppoeSubProfString
- TmnxLocUsrDbDhcpEntry
    - tmnxLocUsrDbDhcpSlaProfString
    - tmnxLocUsrDbDhcpSubProfString
- TmnxSubPppModEntry
    - tmnxSubPppModSubProfStr
    - tmnxSubPppModSlaProfStr
- TmnxSubIpoeModEntry
    - tmnxSubIpoeModSubProfStr
    - tmnxSubIpoeModSlaProfStr
- TmnxSubSlaacModifyEntry
    - tmnxSubSlaacModifySubProfString
    - tmnxSubSlaacModifySlaProfString
- SvcDhcpLeaseModifyEntry
    - svcDhcpLeaseModifySubProfile
    - svcDhcpLeaseModifySlaProfile
- svcDhcpLeaseTable
    - svcDhcpLeaseSubProfString
    - svcDhcpLeaseSlaProfString
- svcArpHostTable
    - svcArpHostSubProfString
    - svcArpHostSlaProfString
- tmnxSubSlaacTable
    - tmnxSubSlaacSubProfString

        - tmnxSubSlaacSlaProfString
    - tmnxSubIpoeTable
        - tmnxSubIpoeSubProfString
        - tmnxSubIpoeSlaProfString
    - tmnxSubPppTable
        - tmnxSubPppSubProfString
        - tmnxSubPppSlaProfString

CLI show commands:

```
show service id 10 ipoe session detail
show service id 10 pppoe session detail
show service id 10 slaac host detail
show service id 10 dhcp lease-state detail
show service id 10 arp-host detail
show service id 10 dhcp6 lease-state detail
```

CLI tools commands:

```
tools perform subscriber-mgmt edit-ipoe-session
tools perform subscriber-mgmt edit-lease-state
tools perform subscriber-mgmt edit-ppp-session
tools perform subscriber-mgmt edit-slaac-host
```

8.1.7 Release 19.7.R2

In Release 19.7.R2, no commands are changed or removed in classic interfaces.

8.1.8 Release 19.7.R1

8.1.8.1 Ingress Multicast Path Management

- In Release 19.7.R1, an additional capacity of 17000 has been added to the **configure mcast-management chassis-level per-mcast-plane-capacity total-capacity** command for the capacity of IMPM paths and planes with FP4 hardware in a 7750 SR-7-B/12-B.

8.1.8.2 MPLS

- In Release 19.7.R1, two new MPLS commands are introduced for a SR-TE LSP and a RSVP-TE LSP. These commands are intended to replace two existing commands which will be removed in a subsequent release.

  The co-existence of the existing commands and new commands is accommodated in Release 19.7.R1, so either can be configured. However, the CLI information and the MIB table used is that of the new command. Table 18 and Table 19 show the effect of each user-entered command on each LSP type in Release 19.7.R1. The same behavior described for the existing command is also applied when upgrading a configuration to Release 19.7.R1.

Table 18 **RSVP-TE Commands**

		LSP		**LSP Template**	
New Command	**Existing Command**	**Action**	**CLI Information/ MIB**	**Action**	**CLI Information/ MIB**
--	**no cspf**	Yes	**no path-computation-method** (IGP LSP)	None (cspf is mandatory and default)	**path-computation-method local-cspf**; **metric-type igp**
--	**cspf**	Yes	**path-computation-method local-cspf**; **metric-type igp**	Yes	**path-computation-method local-cspf**; **metric-type igp**
--	**cspf use-te-metric**	Yes	**path-computation-method local-cspf**; **metric-type te**	Yes	**path-computation-method local-cspf**; **metric-type te**
no path-computation-method	--	Yes	**no path-computation-method** (IGP LSP)	None (cspf is mandatory and default)	**path-computation-method local-cspf**; **metric-type igp**
path-computation-method local-cspf	--	Yes	**path-computation-method local-cspf**; **metric-type igp**	Yes	**path-computation-method local-cspf**; **metric-type igp**

Table 18 **RSVP-TE Commands (Continued)**

		LSP		**LSP Template**	
New Command	**Existing Command**	**Action**	**CLI Information/ MIB**	**Action**	**CLI Information/ MIB**
path-computation-method pce	--	Yes	**path-computation-method pce**	N/A (blocked)	None
no metric-type	--	Yes	**metric-type igp**	Yes	**metric-type igp**
metric-type igp	--	Yes	**metric-type igp**	Yes	**metric-type igp**
metric-type te	--	Yes	**metric-type te**	Yes	**metric-type igp**
--	**no pce-computation**	Yes	Value of **path-computation-method**	N/A (blocked)	None
--	**pce-computation**	Yes	**path-computation-method pce**	N/A (blocked)	None

Table 19 **SR-TE Commands**

		LSP		**LSP Template**	
New Command	**Existing Command**	**Action**	**CLI Information/ MIB**	**Action**	**CLI Information/ MIB**
--	**no cspf**	Yes	**no path-computation-method** (ip-to-label method)	None (cspf is mandatory and default)	**no path-computation-method** (ip-to-label method)
--	**cspf**	Yes	**no path-computation-method** (ip-to-label method)	Yes	**no path-computation-method** (ip-to-label method)
--	**cspf use-te-metric**	Yes	**no path-computation-method** (ip-to-label method)	Yes	**no path-computation-method** (ip-to-label method)

Table 19 **SR-TE Commands (Continued)**

		LSP		**LSP Template**	
New Command	**Existing Command**	**Action**	**CLI Information/ MIB**	**Action**	**CLI Information/ MIB**
no path-computation-method	--	Yes	**no path-computation-method** (ip-to-label method)	Yes	**no path-computation-method** (ip-to-label method)
path-computation-method local-cspf	--	Yes	**path-computation-method local-cspf**	Yes	**path-computation-method local-cspf**
path-computation-method pce	--	Yes	**path-computation-method pce**	N/A (blocked)	None
no metric-type	--	Yes	**metric-type igp**	Yes	**metric-type igp**
metric-type igp	--	Yes	**metric-type igp**	Yes	**metric-type igp**
metric-type te	--	Yes	**metric-type te**	Yes	**metric-type te**
--	**no pce-computation**	Yes	Value of **path-computation-method**	N/A (blocked)	None
--	**pce-computation**	Yes	**path-computation-method pce**	N/A (blocked)	None

8.1.8.3 QoS

- In Release 19.7.R1, the **copy** command within the **config**>**qos** context has been modified to become a branch in the CLI and the policies to be copied are configured as sub-commands. The overall syntax has not changed; however, it is now mandatory that the first parameter is the name of the policy to be copied, as reflected in the CLI help. Entering only the **config**>**qos**>**copy** command causes the CLI to enter the **copy** context which requires an **exit** command to exit. [320018]

8.1.8.4 NAT

- In Release 19.7.R1, the **tools perform nat recover-l2aw-bypass** *mda-id* command is changed to **tools perform nat recover-l2aw-bypass mda** *mda-id*.

8.1.8.5 WLAN-GW

- In Release 19.7.R1, the following commands under the **configure service template epipe-sap-template** *name* context no longer allow references to filters or QoS policies that are not present on the system:
  - **egress filter ip**
  - **egress filter ipv6**
  - **egress filter mac**
  - **egress qos**
  - **ingress filter ip**
  - **ingress filter ipv6**
  - **ingress filter mac**
  - **ingress qos**

  The existing configurations with such references will be upgraded gracefully and have these references automatically removed. References to valid filters or QoS policies are not impacted.

8.1.9 Release 19.5.R2

In Release 19.5.R2, no commands are changed or removed in classic interfaces.

8.1.10 Release 19.5.R1

8.1.10.1 System

- In Release 19.5.R1, the following **psi-tti** commands in the **config**>**port**>**otu** context have been removed:

```
configure port otu psi-tti
configure port otu psi-tti expected auto-generated
```

```
     configure port otu psi-tti expected bytes <byte-string> [<byte-
string>...(up to
          64 byte-strings max, 255 bytes max)]
     configure port otu psi-tti expected string <identifier>
     configure port otu psi-tti expected use-rx
     configure port otu psi-tti mismatch-reaction {squelch-rx}
     configure port otu psi-tti tx
     configure port otu psi-tti tx auto-generated
     configure port otu psi-tti tx bytes <byte-string> [<byte-string>...(up to 64
          byte-strings max, 255 bytes max)]
     configure port otu psi-tti tx string <identifier>
```

Before upgrading a node to a release where these commands have been removed, ensure that all OTU ports connecting to the node set the mismatch reaction to disable to avoid any impact to the traffic flow. [300425]

- In Release 19.5.R1, the commands associated with configuring 7450 ESS mixed-mode chassis behavior have been removed. 7450 ESS mixed-mode chassis behavior is now always enabled.
  - **config**>**card**>**capability**
  - **config**>**system**>**mixed-mode**
  - **tools**>**perform**>**system**>**mixed-mode-upgrade**
- In Release 19.5.R1, the **cflowd rate** command is changed in classic interfaces.

  Command prior to Release 19.5.R1:
  - **config**>**cflowd**>**rate** *sample-rate*

  Commands in Release 19.5.R1:
  - **config**>**cflowd**>**sample-profile** *profile-id* [**create**]
  - **config**>**cflowd**>**sample-profile sample-rate** [*1..10000*] (Default - 1000)
- In Release 19.5.R1, the following **mcpath** log events have been replaced.

  Log events prior to Release 19.5.R1:

  ```
  tmnxMcPathSrcGrpBlkHole (2001)
  tmnxMcPathSrcGrpBlkHoleClear (2002)
  tmnxMcPathAvailBwLimitReached (2003)
  tmnxMcPathAvailBwValWithinRange (2004)
  ```

  Log events in Release 19.5.R1:

  ```
  tmnxMcPathSrcGrpBlackHole (2005)
  tmnxMcPathSrcGrpBlackHoleCleared (2006)
  tmnxMcPathAvailBwLimitExceeded (2007)
  tmnxMcPathAvailBwLimitCleared (2008)
  ```

  The new log events must be used in the following contexts:

  ```
  configure log event-control mcpath
  configure log event-trigger event mcpath
  ```

Any configuration that references the old **mcpath** events 2001 to 2004 must be removed from the configuration before an upgrade to Release 19.5.R1. If configuration of the **mcpath** events 2001 to 2004 exists in the saved classic CLI configuration file during an upgrade, then the configuration file will fail to load after the upgrade.

8.1.10.2 NETCONF

- The **admin**>**system**>**candidate**>**discard-changes** command is removed. Several other commands to discard changes in the candidate using model-driven interfaces were introduced in Release 16.0. [309473]

8.1.10.3 Filters

- In SR OS Release 19.5.R1, the commands (and all sub commands) of **config**>**filter**>**redirect-policy**>**destination**>**snmp-test** and **config**>**filter**>**redirect-policy**>**destination**>**url-test** are removed:

```
configure filter redirect-policy destination snmp-test
configure filter redirect-policy destination snmp-test drop-count
configure filter redirect-policy destination snmp-test interval
configure filter redirect-policy destination snmp-test oid
configure filter redirect-policy destination snmp-test return-value
configure filter redirect-policy destination snmp-test timeout

configure filter redirect-policy destination url-test
configure filter redirect-policy destination url-test drop-count
configure filter redirect-policy destination url-test interval
configure filter redirect-policy destination url-test return-code
configure filter redirect-policy destination url-test timeout
configure filter redirect-policy destination url-test url
```

8.1.10.4 Routing

- In Release 19.5.R1, the following **bgp multipath** commands are changed.

  Command prior to Release 19.5.R1:

```
configure router bgp multipath <max-paths> [ebgp <ebgp-max-paths>] [ibgp <ibgp-max-
     paths>] [restrict {same-neighbor-as|exact-as-path}]
configure service vprn bgp multipath <max-paths> [ebgp <ebgp-max-
paths>] [ibgp <ibgp-
     max-paths>] [restrict {same-neighbor-as|exact-as-path}]
```

  Command in Release 19.5.R1:

```
configure router bgp multi-path maximum-paths <max-paths> [ebgp <ebgp-max-
paths>]      [ibgp <ibgp-max-paths>] [restrict {same-neighbor-as|exact-as-path}]
configure service vprn bgp multi-path maximum-paths <max-paths> [ebgp <ebgp-max-
     paths>] [ibgp <ibgp-max-paths>] [restrict {same-neighbor-as|exact-as-path}]
```

8.1.10.5 QoS

- In release 19.5.R1, the configuration of network ingress pools and queues has moved from under the **card mda network ingress** CLI branch to the **card fp ingress network** branch to allow per-FP configuration on multi-FP XMAs.

  Commands prior to Release 19.5.R1:

```
configure
  card <slot-number>
    mda <mda-slot>
      network
        ingress
         [no] pool
            amber-alarm-threshold <percentage>
            no amber-alarm-threshold
            no red-alarm-threshold
            red-alarm-threshold <percentage>
            no resv-cbs
            resv-cbs <percent-or-default>
            resv-cbs <percent-or-default> amber-alarm-action step <percent>
                                          max <[1..100]>
            no slope-policy             slope-policy <name>
          queue-policy <name>
```

  Commands in Release 19.5.R1:

```
configure
  card <slot-number>
    fp <fp-number>
      ingress
         network
           [no] pool
            amber-alarm-threshold <percentage>
            no amber-alarm-threshold
            no red-alarm-threshold
            red-alarm-threshold <percentage>
            no resv-cbs
            resv-cbs <percent-or-default>
            resv-cbs <percent-or-default> amber-alarm-action step <percent>
                                          max <[1..100]>
            no slope-policy
            slope-policy <name>
          queue-policy <name>
```

8.2 Changed or Removed Commands in Model-driven Interfaces

8.2.1 Release 19.10.R6

In Release 19.10.R6, no commands are changed or removed in model-driven interfaces.

8.2.2 Release 19.10.R5

8.2.2.1 Routing

- In Release 19.10.R5, the allowed range in the **service ies|vprn interface ipv4|ipv6 neighbor-discovery host-route populate** {**static** | **dynamic** | **evpn**} **route-tag** *number* context has been modified from 0 to 255 to 1 to 255. The value of 0 was incorrectly allowed in previous releases. [344313, 344791]

8.2.3 Release 19.10.R4

In Release 19.10.R4, no commands are changed or removed in model-driven interfaces.

8.2.4 Release 19.10.R3

8.2.4.1 System

- YANG “must” statements were introduced as a limited support feature in 19.10.R1. They have now been removed from the YANG models.

8.2.4.2 MPLS

- The following issues are present in the YANG state output for MPLS: [341366]
  - An IPv4/IPv6 operational down reason (MPLS interface and MPLS base context) is not populated if the MPLSv4/v6 (MPLS base context) and/or a MPLSv4/v6 interface (MPLS interface context) operational state becomes operational down.
  - A record label against the first egress hop for an RSVP-TE LSP is incorrectly set to a value of 4294967295.
  - An actual route hop list and computed hop list for a localCspf LSP incorrectly shows the hop as loose with “isLoose true”.

8.2.4.3 Subscriber Management

- When the **user-name-format** is either **dhcp-vendor** or **circuit-id**, then a **mac-format** of “ieee” is not supported. During an upgrade a code sets the **mac-format** value to its default value “alu”. [338901]

8.2.5 Release 19.10.R2

8.2.5.1 MD-CLI

- The MD-CLI **ping** command **wait** parameter is renamed to **timeout** to be consistent with the classic CLI and the ASAA **icmp-ping timeout** command. [337274]

8.2.6 Release 19.10.R1

8.2.6.1 System

- In Release 19.10.R1, the default value of the **configure system management-interface yang-modules base-r13-modules** command is changed to **false**.

- In Release 19.10.R1, the default value of the **configure system management-interface netconf capabilities writable-running** command is changed to **false**.
- Prior to Release 19.10.R1, the type used for **configure card slot-number** was a union of typedef iom-card-slot and typedef cpm-card-slot. However, the **card** commands are only used for IOMs and XCMs, and not for CPMs. In Release 19.10.R1, the type has been changed from the union to just typedef iom-card-slot. [328021]
- Prior to Release 19.10.R1, a number of incorrect default statements existed in the YANG models for configuration groups. In Release 19.10.R1, these default statements have been removed. [328828]
- The type used for the YANG leaf nokia **state port transceiver vendor-oui** has changed from a uint32 to a new string type representing the 3 octets in hexadecimal format. [330980]

8.2.6.2 Routing

- In Release 19.10.R1, the following **state router mpls** leafs have been replaced. [328898]

  Leafs prior to Release 19.10.R1:

  - **state router mpls oper-state**
  - **state router mpls interface oper-state**

  Leafs in Release 19.10.R1:

  - **state router mpls ipv4-oper-state**
  - **state router mpls interface ipv4-oper-state**
  - **state router mpls ipv6-oper-state**
  - **state router mpls interface ipv6-oper-state**
- In Release 19.10.R1, the type of the following leafs in the Nokia YANG model is changed from uint32 into uint64 to support higher values. [328605]
  - **state router rsvp interface bandwidth-info reserved-bw**
  - **state router rsvp interface total-bw**
- In Release 19.10.R1, the range of the **mpls lsp secondary path-preference** command is corrected to 1..255. [329863]

8.2.6.3 MPLS

- In Release 19.10.R1, the **configure router mpls sr-te-resignal-timer** *minutes* command is replaced with the **configure router mpls sr-te-resignal resignal-timer** *minutes* command under a new node. [328682]

8.2.6.4 Subscriber Management

- In Release 19.10.R1, the YANG pattern of the **bd-mac-prefix** leaf has been changed to only permit lower case hexadecimal numbers in the following commands. [332578]
  - **service ies subscriber-interface group-interface wlan-gw vlan-range vrgw lanext bd-mac-prefix**
  - **service vprn subscriber-interface group-interface wlan-gw vlan-range vrgw lanext bd-mac-prefix**

8.2.7 Release 19.7.R2

In Release 19.7.R2, no commands are changed or removed in model-driven interfaces.

8.2.8 Release 19.7.R1

8.2.8.1 Ingress Multicast Path Management

- In Release 19.7.R1, an additional capacity of 17000 has been added to the **configure multicast-management chassis-level per-mcast-plane-capacity total-capacity** command for the capacity of IMPM paths and planes with FP4 hardware in a 7750 SR-7-B/12-B chassis.

8.2.8.2 MPLS

- In Release 19.7.R1, two new MPLS commands are introduced for a SR-TE LSP and a RSVP-TE LSP. These commands are intended to replace two existing commands which will be removed in a subsequent release. See Table 18 and Table 19 for more information.

8.2.9 Release 19.5.R2

In Release 19.5.R2, no commands are changed or removed in model-driven interfaces.

8.2.10 Release 19.5.R1

8.2.10.1 System

- In Release 19.5.R1, the following **mcpath** log events have been replaced.

  Log events prior to Release 19.5.R1:

  ```
  tmnxMcPathSrcGrpBlkHole (2001)
  tmnxMcPathSrcGrpBlkHoleClear (2002)
  tmnxMcPathAvailBwLimitReached (2003)
  tmnxMcPathAvailBwValWithinRange (2004)
  ```

  Log events in Release 19.5.R1:

  ```
  tmnxMcPathSrcGrpBlackHole (2005)
  tmnxMcPathSrcGrpBlackHoleCleared (2006)
  tmnxMcPathAvailBwLimitExceeded (2007)
  tmnxMcPathAvailBwLimitCleared (2008)
  ```

  The new log events must be used in the following contexts:

  ```
  configure log log-events mcpath
  configure log event-trigger mcpath
  ```

  Any configuration that references the old **mcpath** events 2001 to 2004 must be removed from the configuration before an upgrade to Release 19.5.R1. If configuration of the **mcpath** events 2001 to 2004 exists in the saved model-driven configuration file during an upgrade, then the configuration file will fail to load after the upgrade.

- Prior to Releases 19.5.R1, an explicit **mep mac-address** configuration of *00:00:00:00:00:00* was allowed for the remote MEP (remote MAC *00:00:00:00:00:00*) and the local MEP (**mac-address** *00:00:00:00:00:00*). In Releases 19.5.R1 or higher, the explicit configuration is not allowed for remote-MEP or local-MEP configurations. Attempts to load configurations with these commands in subsequent releases will fail. Remove this explicit configuration if it exists before upgrading to Release 19.5.R1 or higher to avoid the issue. [308624]
- In Release 19.5.R1, two log event names in MD Interfaces have changed to align with their names in the classic CLI. Logger event 'started' has changed to 'STARTED', and PCAP event 'tmnxPcapReadWriteFailure' has changed to 'tmnxPcapBufferReadWriteFailure'. [315787]
- In Release 19.5.R1, the default of the **configure system login-control login-banner** leaf has been changed from 'true' to 'false'.
- In Release 19.5.R1, the **cflowd rate** command is changed in model-driven interfaces.

  Command prior to Release 19.5.R1:

  − **configure cflowd rate** *sample-rate*

  Commands in Release 19.5.R1:

  − **configure cflowd sample-profile** *profile-id*

  − **configure cflowd sample-profile sample-rate** [*1..10000*] (Default - 1000)
- The YANG definition of the following two state leaves has been changed from a union type that included a string to a number with two decimal places. [305309]

```
state system cpu cpu-usage
state system cpu capacity-usage
```

8.2.10.2 Routing

- L2TPv3 configuration no longer accepts a local interface address to be configured as an L2TPv3 local-address. This type of configuration was accepted in prior releases but was not supported. [314657]

8.2.10.3 QoS

- The default value specified for the rates under a level of a **port-scheduler-policy** QoS override applied at the egress of an Ethernet port has been removed. The absence of the leaf means that the system will not override these rates. [302444]

- The reference-bandwidth typedef in the Nokia YANG model was changed from uint32 into uint64 to support higher values. [305557]
- The YANG definition of the forwarding class in the following commands has been changed from a leafref to a string length of 32 characters; however, the system continues to only accept the forwarding class name keywords, with an optional sub-class. [306649]

```
configure qos sap-ingress default-fc
configure qos sap-ingress mac-criteria entry action fc
configure qos sap-ingress ip-criteria entry action fc
configure qos sap-ingress ipv6-criteria entry action fc
configure qos sap-ingress dscp fc
configure qos sap-ingress dot1p fc
configure qos sap-ingress prec fc
configure qos sap-ingress lsp-exp f
```

- The YANG definition of the policer identifier in the following commands has been changed from a leafref to an integer length of 32 bits; however, the system continues to only accept the allowed policer identifier range, specifically, 1 to 63. [306649]

```
configure qos sap-ingress fc queue-policer-mapping policer policer
configure qos sap-ingress fc multicast-queue-policer-mapping multicast-
    policer multicast-policer
configure qos sap-ingress fc broadcast-queue-policer-mapping broadcast-
    policer broadcast-policer
configure qos sap-ingress fc unknown-queue-policer-mapping unknown-
    policer unknown-policer
```

- The YANG definition of the queue identifier in the following commands has been changed from a leafref to an integer length of 32 bits; however, the system continues to only accept the allowed queue identifier range, specifically, 1 to 8. [306649]

```
configure qos sap-egress ip-criteria entry action queue
configure qos sap-egress ipv6-criteria entry action queue
```

- The YANG definition of the arbiter and scheduler name in the following commands has been changed from a leafref to a string length of up to 255 characters; however, the system continues to only accept the allowed arbiter and scheduler name length, specifically, 1 to 32 characters. [306649]

```
configure qos policer-control-policy tier arbiter arbiter-parent arbiter-name
configure qos scheduler-policy tier scheduler parent-mapping scheduler-
    parent scheduler-parent scheduler-name
```

- The YANG definition of the **id** leaf (**log-id** key leaf) in the following commands has been changed from a leafref to an integer length of 32 bits. [306649]

```
configure service vprn log log-id id
state service vprn log log-id id
```

- In Release 19.5.R1, the configuration of network ingress pools and queues has moved from under the **card mda network ingress** CLI branch to the **card fp ingress network** branch to allow per-FP configuration on multi-FP XMAs.

  Commands prior to Release 19.5.R1:

```
configure
    card <iom-card-slot>
        mda <number>
            network
                ingress
                    pool <string>
                        amber-alarm-threshold <number>
                        apply-groups <reference>
                        red-alarm-threshold <number>
                        resv-cbs
                            amber-alarm-action
                                max <number>
                                step <number>
                            cbs <number>
                        slope-policy <reference>
                    queue-policy <reference>
```

  Commands in Release 19.5.R1:

```
configure
    card <iom-card-slot>
        fp <number>
            ingress
                network
                    pool <string>
                        amber-alarm-threshold <number>
                        apply-groups <reference>
                        red-alarm-threshold <number>
                        resv-cbs
                            amber-alarm-action
                                max <number>
                                step <number>
                            cbs <number>
                        slope-policy <reference>
                    queue-policy <reference>
```

9 Software Upgrade Procedures

The following sections contain information for upgrading to the Release 19.10.R6 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.
- ISSU Upgrade Procedure

  Procedure for performing an ISSU to Release 19.*x* including information on applicability of ISSU for earlier versions.
- 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 19.10.R6.

Note:

- When using the standard software upgrade procedure to change the major release running on a node, the release date of the target load (for example,16.0.Ry) should be at least 90 days later than the release date of the current load (for example, 15.0.Rx).
- 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 when **configure system management-interface configuration-mode** is set to **classic** (the default setting):

- Obsoleted commands are not flagged as errors upon reading a configuration file with obsoleted 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 19.10.R3 or Higher

- Release 19.10.R3 introduces a non-mandatory firmware upgrade for the me10 -10gb-sfp+ and me6-10gb-sfp+ MDAs. Soft Reset is supported for these cards during an ISSU from an image prior to Release 19.10.R3 to a Release 19.10.R3 or later image. A hard reset of the card has to be performed independently after the ISSU to upgrade the firmware. [338476]

9.1.2 Upgrading to Release 19.10.R2 or Higher

- Upon upgrading to Release 19.10.R2 or higher, operators should validate existing IPv6-in-IPv4 or IP-over-GRE-over-IPv4 use-cases and confirm that the system IP address or the GRE termination subnet is used as the outer tunnel destination address. Network operators may have relied unintentionally on CPM forwarding for this traffic by targeting a router interface IP address instead of the system IP address or the GRE termination subnet. This forwarding using a router interface IP address is now blocked and customers can enable **forward-6-in-4** or **forward-ip-over-gre** at the system level as needed. [326304]
- When upgrading to 19.10.R2 or higher, the commands **adv-config-policy**>**child-control**>**bandwidth-distribution**>**above-offered-allowance**>**unconsumed-higher-tier-rate** *percent* and **adv-config-policy**>**child-control**>**bandwidth-distribution**>**above-offered-allowance**>**delta-consumed-higher-tier-rate** *percent* must be removed from the configuration and re-added after the upgrade. [345778]

9.1.3 Upgrading to Release 19.10.R1 or Higher

- Upon upgrading to Release 19.10.R1, a PE that previously sent EVPN-MPLS routes with non-system IPv4 address P1 as a next-hop, may now send the routes with the system IPv4 causing remote nodes to fail next-hop resolution and service disruption. To avoid the service impact, prior to an upgrade, an export policy must be configured with “action accept next-hop P1”. This policy will guarantee that the non-system IPv4 address P1 is still used as next-hop after the upgrade.
- In Release 19.10.R1, the GTP S11 Session forwarding model is modified so that it is not possible to upgrade an older system with live S11 sessions. Attempting an upgrade results in a reboot of the CPM running the newer image. To upgrade, Nokia recommends shutting down all **group-interfaces** of type **gtp** and clearing all S11 sessions using the **clear subscriber-mgmt gtp session all** command. Once the upgrade has completed, all **gtp group-interfaces** can be re-enabled and sessions will be re-established.
- Release 19.10.R1 introduces a mandatory firmware upgrade on me2-100gb-cfp4 and me2-100gb-qsfp28 MDAs. [251111]
- In Release 19.10.R1, the CLI commands used to control SR OS NETCONF use of the writable-running capability and the ALU base-r13 YANG modules have changed to be disabled/false by default. This is to prevent any accidental usage of the ALU base-r13 modules. [321769]
- In Release 19.10.R1, the configuration of an interface name as IPv4 source address for SR OS applications is made consistent between classic, mixed and model-driven management interface configuration modes. The specified interface must have a valid primary IPv4 address configured. Previously, invalid configurations were accepted but not active.

  Supported interface names as source-address in the Base router instance (**configure system security source-address application** *ipv4-app interface-name*):

```
configure router interface <name>
configure router interface <name> control-tunnel
configure router interface <name> gmpls-loopback
configure service ies interface <name>
configure service ies interface <name> tunnel
configure service ies redundant-interface <name>
configure service ies subscriber-interface <name>
```

  Supported interface names as source-address in a VPRN instance (**configure service vprn** *id* **source-address application** *ipv4-app interface-name*):

```
configure service vprn interface <name>
configure service vprn interface <name> tunnel
configure service vprn network-interface <name>
configure service vprn redundant-interface <name>
configure service vprn subscriber-interface <name>
```

Nokia recommends scrubbing the configuration prior to upgrading a router from an older SR OS release and make the necessary corrections. When upgrading, invalid interface names used as source address will be ignored and removed from the configuration. [321824]

- On VSR, for retail subscriber-interfaces, **dhcp lease-populate** and **ipoe-session session-limit** are now limited to 131071. Saved configuration files that have larger values must be manually updated before execution. [323477]
- The 7950 XRS 12-port Universal QSFP-DD XMA (x12-400g-qsfpdd) was introduced in Release 19.7.R1, but some of the MDA license levels it uses are incorrect. This XMA currently incorrectly reuses a previously defined set of FP4 MDA license levels that are not appropriate for this 7950 XMA. If this 7950 XMA is deployed in Release 19.7.R1 or 19.7.R2 using any of the MDA license levels **cr2400g+**, **er2400g+**, or **he2400g+**, then when this configuration is processed by a later SR OS release containing the correct MDA license levels, the configuration would fail to execute. A manual modification of the XMA configuration will be required as part of the upgrade process to the later release that resolves this issue, containing the correct MDA license levels. [331260]
- On 7750 SR-7/12 systems using SFM6 fabric (m-sfm6-7/12), changes have been made to improve inter-card communication. These changes mandate a standard software upgrade rather than an ISSU operation for the upgrades from Releases 16.0.R7, 16.0.R8, or 16.0.R9 to Release 19.10.R1 or higher.

  If ISSU procedures are used for any of these upgrade paths, please contact Nokia support for assistance. [322406]
- Prior to Release 19.10.R1, the **configure service** {**vprn**|**ies**} *service-name* **subscriber-interface** *interface-name* **group-interface** *group-interface-name* **wlan-gw vlan-range** *range* **vrgw lanext bd-mac-prefix** leaf allowed configuration of MAC prefixes with uppercase characters. In Release 19.10.R1, this has been blocked. Prior to upgrading to Release 19.10.R1 in model-driven configuration mode, any uppercase prefix must be converted to lower-case. [332578]
- It is recommended not to use soft reset on FP4 based line cards in Release 19.10.R1, but soft reset for ISSU from Release 19.10.R1 to a future release can be executed safely. Hard reset (**clear card** *slot*) in Release 19.10.R1 can be used safely as well and if a soft reset was done inadvertently, a hard reset of the MDA/XMA (**clear mda** *slot*) can be executed to recover. [335847]
- In Release 19.10.R1, SR OS allows the PTP profile to be set to g8275dot2-2016 via SNMP on 7750 SR, 7450 ESS, and 7950 XRS platforms; however this profile is not supported. If this PTP profile is configured, it must be changed to one of the following supported profiles prior to upgrading to Release 19.7.R2 or higher: g8265dot1-2010, ieee1588-2008, or g8275dot1-2014. [338174]

9.1.4 Upgrading to Release 19.7.R1 or Higher

- Prior to Release 19.7.R1, it was possible to have a connector on the es64-10gb-sfpp+4-100gb-qsfp28 satellite administratively disabled, but the ports within that connector (the satellite uplink ports) would remain operationally enabled and passing traffic. When upgrading from an earlier release to Release 19.7.R1 or higher, then there is a risk that the upgrade shall cause the satellite uplink ports to change to operationally down. For connectors on satellites that are administratively disabled, and where the ports of that connector are administratively enabled, the connector should be changed to administratively enabled before the upgrade to avoid a change in traffic flow after the upgrade. [322687]
- In Release 19.7.R1, the SR OS gRPC gNMI interface has been changed to OpenConfig gnmi.proto version 0.7.0. Prior to upgrading to Release 19.7.R1, clients and collectors should be updated. However, all 0.7.0 upgrades are backward compatible; thus using gNMI clients with gnmi v.0.4.0 will not cause any instability.

9.1.5 Upgrading to Release 19.7.R1

- The format of **hash2** strings in encrypted values of configuration elements that are part of configuration groups in releases prior to 19.7.R1 are incompatible with the format used in Release 19.7.R1. Configuration files generated in SR OS Releases 16.0 or 19.5 containing **hash2** strings in configuration groups will fail to load on a router running SR OS Release 19.7.R1. [329387]

9.1.6 Upgrading to Release 19.5.R1 or Higher

- The OTU **psi-tti** functionality has been removed. During a network upgrade, if there is an OTU connection between nodes, and one node is upgraded to Release 19.5.R1 or higher, and the other node remains on an SR OS release prior to Release 19.5.R1, and the OTU connection has **psi-tti** configuration, then a configuration mismatch may result. Nokia recommends disabling the **psi-tti** configuration on the OTU ports on both ends of the connection before the upgrade to avoid any issues. At a minimum, **psi-tti mismatch-reaction squelch-rx** should be disabled to ensure no loss of traffic after a one-sided upgrade. [300425]

- In Release 19.5.R1, the configuration of buffer pool information and network queues has been moved from the **card mda network ingress** CLI tree to the **card fp ingress network** CLI tree. When a configuration file containing the pre-Release 19.5.R1 (MDA) commands is executed in the classic CLI, or loaded in the MD-CLI with a **load full-replace**, after rebooting to Release 19.5.R1 or higher, the pre-Release 19.5.R1 (MDA) commands are automatically converted to the equivalent Release 19.5.R1 or higher (FP) commands. When upgrading to Release 19.5.R1 or higher in model-driven configuration mode, and the existing MDA commands are part of a configuration group applied (using the **apply-groups** command) in the pre-Release 19.5.R1 configuration file that is loaded with a **load full-replace**, the commands prior to Release 19.5.R1 (see Changed or Removed Commands for a list of the commands) are:

  1. Expanded into the base configuration
  2. Removed from the configuration group
  3. Upgraded to the commands in Release 19.5.R1 or higher (see the QoS section of Changed or Removed Commands) within the base configuration

  If the configuration group is not applied, only Step 2 is performed. Other commands in the configuration group are unaffected.

  To continue to use configuration groups for these commands (now under the FP), after the upgrade is completed, create an appropriate configuration group with the related commands, apply this configuration group above the commands, and remove the commands from the base configuration.

- In Release 16.0, SR OS systems that have Pay-As-You-Grow Hardware License upgrades applied, require special handling to upgrade to Release 19.5. For a system that has a license file specified in the BOF, and has upgrades applied to FP4 IOMs and XMAs, then the license file must be updated to include a key for Release 19.5 prior to performing the upgrade.

  This update is performed by editing the Network Function in Nokia Centralized License Manager (CLM) to update its "Version" to 19, then using the “Generate or Deploy a new license key” to re-issue the license file to the node. The **show**>**system**>**license available-licenses** command can be used to check the releases covered by the installed license file.

- When upgrading to Release 19.5.R1 in model-driven configuration mode, all configuration related to **mcpath** log events must be removed prior to the upgrade. The following **mcpath** log events have been replaced:

  ```
  tmnxMcPathSrcGrpBlkHole (2001)
  tmnxMcPathSrcGrpBlkHoleClear (2002)
  tmnxMcPathAvailBwLimitReached (2003)
  tmnxMcPathAvailBwValWithinRange (2004)
  ```

  The following are the new log events in Release 19.5.R1:

  ```
  tmnxMcPathSrcGrpBlackHole (2005)
  tmnxMcPathSrcGrpBlackHoleCleared (2006)
  tmnxMcPathAvailBwLimitExceeded (2007)
  ```

```
tmnxMcPathAvailBwLimitCleared (2008)
```

The new log events must be used in the following contexts:

```
configure log log-events mcpath
configure log event-trigger mcpath
```

Any configuration that references the old **mcpath** events 2001 to 2004 must be removed from the configuration before an upgrade to Release 19.5.R1. If configuration of the **mcpath** events 2001 to 2004 exists in the saved model-driven configuration file during an upgrade, then the configuration file will fail to load after the upgrade.

- The YANG module alu-netconf-deviations-r13.yang has been removed and replaced with nokia-ietf-netconf-deviations.yang in Release 19.5.R1.
- Starting in Release 19.5.R1, the bootup processing of the 7950 XRS-40 has changed. In previous releases, the extension chassis waits for the active CPM of the master chassis to transfer the boot files (boot.ldr, BOF, and image) during the boot process. Beginning in Release 19.5.R1, the CPMs of the extension chassis load their image files either from their local compact flash, or using FTP from a URL reachable by their individual management ports. To support this change, the boot.ldr, BOF, and image files must be copied to the extension chassis CPMs before the reboot.

  On a master chassis running any load of Releases 15.0, 15.1, or 16.0, the **admin redundancy synchronize boot-env** command copies the boot files onto the CPMs of the extension chassis.

  For a master chassis running a load earlier than Release 15.0, or if a manual copy is desired, and where the image and configuration files are to be stored locally on the compact flash, then the following commands can be used where TiMOS-*m.n.Yz* is the directory containing the images for the desired software release. If secondary and tertiary locations are used for images, then these should also be copied to CPM C and D.

```
file copy cf3:/boot.ldr cf3-c:/boot.ldr
file copy cf3:/bof.cfg cf3-c:/bof.cfg
file md cf3-c:/TiMOS-m.n.Yz
file copy cf3:/TiMOS-m.n.Yz/* cf3-c:/TiMOS-m.n.Yz/
file copy cf3:/boot.ldr cf3-d:/boot.ldr
file copy cf3:/bof.cfg cf3-d:/bof.cfg
file md cf3-d:/TiMOS-m.n.Yz
file copy cf3:/TiMOS-m.n.Yz/* cf3-d:/TiMOS-m.n.Yz/
```

For a master chassis running a load earlier than Release 15.0 or if a manual copy is desired, and where the image files are stored remotely, additional steps are required. CPM C and CPM D must be connected to the management network so that they can each independently access the remote locations for the image and configuration files. For this, there must be additional IP addresses for the management port of each of CPM C and CPM D. These addresses need to be added to the BOF. This requires manual editing of the bof.cfg file on the active CPM, and then a copy of this BOF to the standby CPM, CPM C, and CPM D. It also requires a new format for the definition of the standby addresses for the master chassis CPM A and CPM B. In Release 19.5.R1, five IP addresses are required per system. One each for standby/A, standby/B, standby/C, and standby/D and one for the active CPM. To configure boot processing:

1. Use the **file vi** command to edit the bof.cfg file and remove the existing address standby line and add the following lines (as appropriate based on IPv4 and/or IPv6).

```
address a1.b1.c1.d1/n1 standby/a
address p1:q1:r1:s1:t1:u1:v1:w1/m1 standby/a
address a2.b2.c2.d2/n2 standby/b
address p2:q2:r2:s2:t2:u2:v2:w2/m2 standby/b
address a3.b3.c3.d3/n3 standby/c
address p3:q3:r3:s3:t3:u3:v3:w3/m3 standby/c
address a4.b4.c4.d4/n4 standby/d
address p4:q4:r4:s4:t4:u4:v4:w4/m4 standby/d
```

2. Copy the bof.cfg file to the standby CPM using the **admin redundancy synchronize boot-env** command.
3. Copy the bof.cfg and boot.ldr files to the CPMs of the extension chassis using:

```
file copy cf3:/boot.ldr cf3-c:/boot.ldr
file copy cf3:/bof.cfg cf3-c:/bof.cfg
file copy cf3:/boot.ldr cf3-d:/boot.ldr
file copy cf3:/bof.cfg cf3-d:/bof.cf
```

9.1.7 Upgrading to Release 16.0.R7 or Higher

- Release 16.0.R7 introduces a mandatory firmware upgrade on imm40-10g-sfp IMM and x40-10g-sfp XMA cards. A Soft Reset is not supported for these cards during an ISSU from an image prior to Release 16.0.R7; a hard reset will occur instead. [303843]
- The default **trim-mode** for m-sfm5-7 and m-sfm5-12 is now set to **trim-mode-b**. This affects systems where **fabric-speed-a** was being used and where no **trim-mode** has specifically been set by the operator through the **tools perform system set-trim-mode trim-mode-a | b** command. In this case, the default trim-mode **trim-mode-b** is used instead of **trim-mode-a**. In addition, after an ISSU where no **trim-mode** was set by the operator, and where **fabric-speed-a** was used, the system could end up in a state where both trim modes are active. This can be observed using the **tools dump system trim-mode** command. For the new settings to take effect on all IOMs, CPMs and SFMs, clear each SFM in the system using the **clear sfm** *sfm-num* command. To avoid any traffic disruption, the operator should wait for one SFM to be back on-line before attempting to **clear** the other. [313052]

- Release 16.0.R7 introduces a mandatory firmware upgrade for the 7750 SR CPM-2s. During the software upgrade, systems that require the upgrade will automatically update their firmware when they are rebooted as part of the normal software upgrade process. The firmware upgrade will cause a longer reboot time than usual (approximately 10 minutes instead of a few minutes). Ensure the system is not powered down while it is programming the new firmware. [314231]
- Due to the dependency between the uplink configuration and the QSFP28 connector configuration, as of Release 16.0.R6, the 64x10GE+4xQSFP28 satellite does not have a default primary uplink port-mapping configuration. As a result, when configuring the 64x10GE+4xQSFP28 satellite, the port map configuration, including a primary uplink, for all client ports must be explicitly configured.

  Also, configurations saved under Release 16.0.R4 or 16.0.R5, which include the configuration for one or more 64x10GE+4xQSFP28 satellites, may not be compatible with Release 16.0.R6 or higher. To avoid this compatibility issue, before upgrading the 7750 SR host from Release 16.0.R4 or 16.0.R5, a port-map configuration should be made for all client ports of a 64x10GE+4xQSFP28 satellite specifying either a primary and secondary uplink or the primary uplink must be configured to an uplink other than the default uplink. This step is to ensure that a port-map configuration command is saved for all 64x10GE+4xQSFP28 satellite client ports in the Release 16.0.R4 or 16.0.R5 configuration file. Alternatively, ISSU can be used to upgrade a node from Release 16.0.R4 or 16.0.R5 to Release 16.0.R6 or higher without having to add explicit port-map configuration lines for each client port of the 64x10GE+4xQSFP28 satellites. [316512]

9.1.8 Upgrading to Release 16.0.R6 or Higher

- Using '0' as a value for **arp-retry-timer** will result in a failure to execute the configuration when upgrading from an earlier release to Release 16.0.R6 or higher. [270465]
- When a *system-name* is configured with **configure system name** *system-name* with a string longer than 32 characters, the following should be considered. [288438]
  - For a DHCPv6 relay, a group interface can be configured to insert the *system-name* as part of the Interface-ID Option 18 (**dhcp6 option interface-id ascii-tuple**).
  - An SR OS DHCPv6 server drops Relay-Forward messages when the Interface-ID is longer than 64 characters.

- Release 16.0.R6 introduces a mandatory firmware upgrade for the me10-10gb-sfp+ and me6-10gb-sfp+ MDAs. A Soft Reset is not supported for these cards during an ISSU to Release 16.0.R6 or higher; a hard reset will occur instead. [307999]
- Prior to Release 16.0.R6, in MD Interfaces, an explicit **mep mac-address** configuration of *00:00:00:00:00:00* was allowed for the remote MEP (remote MAC *00:00:00:00:00:00*) and the local MEP (**mac-address** *00:00:00:00:00:00)*. The explicit configuration is not allowed for remote-MEP or local-MEP configurations in Releases 16.0.R6 or higher. Attempts to load model-driven configurations with these commands in subsequent releases will fail. Remove this explicit configuration if it exists before upgrading to Release 16.0.R6 or higher to avoid the issue. [308624]
- When upgrading to Release 16.0.R6, the configuration of **scope exclusive** in a SAP egress QoS policy applied to a SAP followed by the configuration of an HS-MDAv2 **packet-byte-offset** queue override on that SAP is no longer permitted. Previously, it was erroneously accepted. This combination should be removed from the configuration. [309762]

9.1.9 Upgrading to Release 16.0.R5 or Higher

- A Major ISSU from a release with enhanced traffic load balancing enabled, to Release 16.0.R5 or higher is now supported for NAT. [308746]
- Upgrading NAT using ISSU, or a reboot from the release with legacy load balancing enabled, to Release 16.0.R5 or higher, must still follow the ‘Preventive Action’ steps outlined in TA 18-0946. If these steps are not followed, the port forwarding entries will still be invalidated during the upgrade, but the ISA will not reset in case of ISSU. When upgrading using reboot, the loading of the configuration file will not be aborted. [308746]
- Upgrading WLAN-GW using ISSU, or a reboot from the release with either legacy or enhanced load balancing enabled to Release 16.0.R5 or higher, must still follow the procedure outlined in the TA 18-0946. If these steps are not followed, the operator will be forced to **shutdown** the WLAN-GW group for ISSU upgrades (and thus follow the TA procedure). When upgrading using reboot, the persistency records will be skipped rather than incorrectly installed. [308746]
- The resolved issue [309931-MA] for the me1-100gb-cfp2 MDA will not become actively resolved after an ISSU from Releases 16.0.R4 or 16.0.R4-1 to Release 16.0.R5 or higher with a Soft Reset of the IOM that contains the me1-100gb-cfp2 MDA. A hard reset of the MDA using the **clear mda** *mda-id* command is required to make the [309931-MA] solution active on the MDA after such an ISSU has been executed.

9.1.10 Upgrading to Release 16.0.R4 or Higher

- A Major ISSU from SR OS releases between 15.0.R4 and 15.0.R7 to Releases 16.0.R4 or higher could result in a hard reset and traffic impact on FP2-based line cards. A workaround is to first perform a minor ISSU to Release 15.0.R8 or higher, followed by a major ISSU. [256308-MA]
- Release 16.0.R4 introduces a mandatory firmware upgrade for me1-100gb-cfp2 MDA. A Soft Reset is not supported for these cards during an ISSU from an image prior to Release 16.0.R4 to a Release 16.0.R4 or higher image; a hard reset will occur instead. [279794]
- When upgrading to Release 16.0.R4 from an earlier release, the configuration combination of BGP-EVPN and a SAP multi-chassis ring-node in the same Epipe service is no longer permitted. Previously, it was erroneously accepted. If both are configured within an Epipe service, one must be removed before upgrading. [295756]
- Release 16.0.R4 introduces a mandatory firmware upgrade for p10-10g-sfp, p6-10g-sfp, p1-100g-cfp, cx20-10g-sfp, and cx2-100g-cfp MDAs/XMAs. A Soft Reset is not supported for these cards during an ISSU from an image prior to Release 16.0.R4 to a Release 16.0.R4 or higher image; a hard reset will occur instead. [299643]
- The combination of the configuration of connection profile VLAN SAPs and MLD snooping in the same VPLS service is no longer permitted (it was incorrectly accepted). If both are configured within a VPLS service, one must be removed before upgrading to Release 16.0.R4. [303479]

9.1.11 Upgrading to Release 16.0.R3 or Higher

- When upgrading SFM cards from FP3-based SFMs to FP4-based SFMs, ensure the system is running with an FP3-based configuration file for the SFMs and XMA/XCMs and the system is running Release 16.0.R3 or higher. To perform the upgrade: [296500]
    1. Execute the **tools perform system set-fabric-speed fabric-speed-c** command. The user will be prompted to confirm that this command is desired.
    2. The node reboots as part of executing the command.
    3. The node loads the configuration file and ignores the configuration lines for the FP3-based SFMs.
    4. Configure the new SFM type for the relevant platform. That is, m-sfm6-12e for the 7750 SR-12e and sfm2-x20s for the 7950 XRS-20/20e.

5. Execute **admin save** to save the new SFM configuration in the configuration file.
6. The existing FP3-based SFMs can be removed and new FP4-based SFMs can be inserted at any point in the procedure after the fabric speed had been set to **fabric-speed-c**.
7. When the standby CPM is online, ensure that the configuration file changes are synchronized with the standby CPM.

9.1.12 Upgrading to Release 16.0.R2 or Higher

- A **gre-tunnel-template** can only support 1023 templates, but 1024 templates can be configured. If an existing Release 16.0.R1 configuration files contain 1024 **gre-tunnel-template**s, the extra template must be removed and saved before upgrading the node to Release 16.0.R2 or higher. [295789]
- When upgrading to Release 16.0.R2 or higher with **configure system management-interface configuration-mode** set to **model-driven**, the configuration (from the previous release) may fail to load. A number of non-backwards-compatible changes were made in MD-CLI configuration commands. Contact your Nokia representative before attempting an upgrade in model-driven configuration-mode.
- When upgrading to Release 16.0.R2 or higher, the **configure multicast-management chassis-level per-mcast-plane-capacity total-capacity 16500** configuration must be removed from the configuration file. [345589]

9.1.13 Upgrading to Release 16.0.R1 or Higher

- When upgrading to Releases 16.0.R1 or higher from an earlier major release, if **bgp-evpn vxlan** is configured in a VPLS service containing a non-default configuration, but no **vpls** *service-id* **vni** *vni* **vxlan instance 1** exists in the same service, the configuration file will fail to execute. In this case, the **bgp-evpn vxlan** configuration must be removed prior to the upgrade.

  Similarly, if a configuration file contains a VPLS service where **bgp-evpn vxlan** and **bgp-evpn mpls** are configured with a non-default configuration (although shut down), the configuration file will fail to execute when upgrading to Release 16.0.R1 or higher. The router will attempt to add **bgp instance 1** to **bgp-evpn mpls**. However, **bgp instance 1** is already associated with VXLAN. As an example, the following configuration in Release 15.0 will fail when upgrading to Release 16.0.R1. [285993]

  ```
  config>service>vpls#
  ```

```
bgp-evpn
  vxlan
      send-evpn-encap
  exit
  mpls
    control-word
    no shutdown
  exit
```

- In Release 16.0.R1, if a BGP IPv6 or multicast-IPv6 route is matched by an export policy rule that tries to change the next-hop to an IPv4 address, the route is treated as though it was rejected by the policy entry. Prior to Release 16.0.R1, the route was considered accepted by the policy entry and the next-hop change was ignored.
- In Release 16.0.R1, if a BGP IPv4 or multicast-IPv4 route is matched by an export policy rule that tries to change the next-hop to an IPv6 address, the route is treated as though it was rejected by the policy entry. Prior to Release 16.0.R1, the route was considered accepted by the policy entry and the next-hop change was ignored.
- In Release 16.0.R1, if a BGP label-IPv4 or VPN-IPv4 route is matched by an export policy rule that tries to change the next-hop to an IPv6 address, and this capability is not supported by the peer or the **advertise-ipv6-next-hops** command is not configured to allow this, then the route is treated as though it was rejected by the policy entry. Prior to Release 16.0.R1, the route was considered accepted by the policy entry and the next-hop change was ignored.
- For 7950 XRS systems, all configuration commands under the **card**>**fp** context for FPs corresponding to an XMA that is not provisioned (which may occur if the configuration file is created using the **admin save detail** command in a release prior to 16.0.R1) must be removed before the upgrade; otherwise, the first related **card**>**fp** command encountered will cause a failure and the processing of the configuration file will terminate. [283741]
- Prior to upgrading to Release 16.0.R1, a *service-name*, *filter-name*, and *policy-name* must be chosen and configured for the following objects if names other than the string-converted ID are desired:
  - all services (**configure service vprn**, **vpls**, **epipe**, etc.)
  - **configure mirror mirror-dest**
  - **configure service pw-template** contexts
  - **configure service customer**
  - **configure filter ip-filter** | **ipv6-filter** | **mac-filter**
  - **configure qos network** |**sap-ingress** | **sap-egress**
  - **configure eth-cfm domain** | **association**

  See Usage Notes for more information.

- The SR OS gRPC telemetry interface in Release 16.0.R1 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.
- When upgrading to Releases 16.0.R1 and higher, the system supports egress IPv6-fragment-match criteria. In prior releases, IPv6-fragment-match criteria were supported on ingress filter; however, the policy could have been configured by the operator on egress. Prior to upgrading to Releases 16.0.R1 or higher, ensure that IPv6 egress filters do not contain fragment match to preserve the same behavior.

9.1.14 Upgrading to Release 15.0.R9 or Higher

- If **single-fiber** has been enabled on an Ethernet satellite port and the port is in **autonegotiate** mode (default for 1G ports), major ISSU and minor ISSU will fail from releases prior to 15.0.R9 to Release 15.0.R9 or higher. A workaround is to remove **single-fiber** before the upgrade. [285066]

9.1.15 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]

9.1.16 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.
- When performing a Major ISSU from Release 14.0 to Release 15.0, the 7210 Ethernet satellite must have its configuration upgraded. To achieve the upgrade, the satellite client ports must be briefly shut down. Although this adds only a brief amount of extra outage time, the far-end port may notice a physical fault and bring down any protocols or services using these ports, resulting in a longer outage. As a workaround, Nokia recommends temporarily reconfiguring the far-end ports to use a higher hold-time until the Major ISSU upgrade is complete.

9.1.17 Upgrading to Release 15.0.R3 or Higher

- After performing an ISSU upgrade from Releases prior to 15.0.R3, existing BGP-VPLS and BGP-VPWS spoke bindings on automatic SDPs will continue to use an LSP type of LDP. If it is required to allow these spoke SDPs to be eligible to also use a BGP LSP type, the spoke SDPs must be re-signaled by performing **a shutdown** then **no shutdown** of **bgp-vpls** and **bgp-vpws**, respectively. [256401]

9.1.18 Upgrading to Release 15.0.R1 or Higher

- Several commands were obsoleted in 15.0.R1. During an ISSU upgrade, the presence of some of these commands in a configuration will cause the standby CPM to fail synchronization when it is rebooted (before the CPM switchover). Refer to the *SR OS 15.0.R1 Software Release Notes* for a complete list of these commands.

- 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 SAS 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, 7950 XRS, and VSR Basic System Configuration Guide* in the “Satellite Software Upgrade Overview” section.
- When performing a standard software upgrade from any Release 14.0 minor release beginning in 14.0.R4 or higher, to Release 15.0.R1 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]

- If upgrading from Releases 13.0 or 14.0 to 15.0.R1, in a VPN route selection, the next-hop IGP cost is taken into account (when no **vrf-import** is present). This may cause an incorrect VPN route to be selected. [262800]

9.1.19 Upgrading to Release 14.0.R7 or Higher

- When performing an ISSU upgrade to Release 14.0.R7 or higher, the MS-ISA and MS- ISA2 cards running isa-bb/isa2-bb applications will continue to operate until their host IOMs upgrade procedure is completed. When the host IOM is upgraded either via the **clear card** *n* **soft hard-reset-unsupported-mdas** or **clear card** *n* commands, the ISA will reboot and load the new image. [246395]
- 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 obsoleted OES commands must be removed prior to an upgrade.
- When performing an ISSU upgrade, from Release 14.0.R4, 14.0.R5, or 14.0.R6 and there is an Optical Extension Shelf (OES) configuration, then the standby CPM will fail to synchronize when it is rebooted (before CPM switchover) and a log event will be generated on the active CPM. These obsoleted OES commands must be removed prior to an ISSU upgrade.

9.1.20 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.

9.1.21 Upgrading to Release 14.0.R5 or Higher

- Release 14.0.R5 introduces a mandatory firmware upgrade for me40-1gb-csfp MDAs. A Soft Reset is not supported for these cards during an ISSU from an image prior to Release 14.0.R5 to a Release 14.0.R5 or higher image; a hard reset will occur instead. [231439]

9.1.22 Upgrading to Release 14.0.R4 or Higher

- During Major ISSU (MISSU) from a release prior to Release 14.0.R4 to Release 14.0.R4 or higher, and especially when the active CPM is still running the older release and the standby CPM is running the new release, no changes should be made to the configuration of the following commands:
  - **advertise-label**
  - **backup-path**
  - **add-paths**
  - **advertise-external**
  - **as-path-ignore**
  - **family**
  - **ebgp-link-bandwidth**
  - **enable-origin-validation**
  - **preference**

  Changes to the above commands could cause the active and standby CPM to become out-of-sync due to the BGP RIB management changes in Release 14.0.R4 and possibly cause the standby CPM to reset.During MISSU, if the **neighbor** context for a BGP session has a **family** command and the session was previously configured with the **advertise-label** command the BGP session will flap during the upgrade.

  Some BGP related commands may not be modified as expected during MISSU. The **prefix-limit**, **ebgp-link-bandwidth**, and **enable-origin-validation** commands all support new **label-ipv4** and **label-ipv6** options starting in Release 14.0.R4. Unfortunately, these new options are not added automatically during MISSU and must be configured after the upgrade is complete. The **add-paths** context supports new **label-ipv4** and **label-ipv6** commands, and while these will be added automatically during MISSU the receive parameter setting may not be as expected in some cases.
- 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
- During a Major ISSU (MISSU) upgrade to Release 14.0, any Time-of-Day (ToD) Suite and Time-Range commands for Filters and QoS obsoleted in Release 14.0.R1 that are in the configuration file will cause the standby CPM to fail synchronization when it is rebooted (before the CPM switchover) and a log event to be raised on the active CPM. These obsoleted commands should be removed prior to a MISSU upgrade.
- 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]
- As of Release 14.0, WLAN-GW uses the IPoE session concept. Consequently, when upgrading from a previous release using major ISSU, only IPv4 states will be kept. All IPv6 states will be lost during the upgrade.

9.1.23 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]

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

  Ignoring the obsoleted 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 obsoleted 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 obsoleted commands, and the configuration would fail to load.

9.1.25 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]

9.1.26 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 (note that ISSU is not supported on the 7750 SR-a platform)
  - 7750 SR-7/12/12e: the CPM5 has new mandatory firmware in Release 13.0.R5 (this does not affect ISSU—CPMs are always rebooted during ISSU)

  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.

9.2 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 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 boot environments must be synchronized if the router has redundant CPMs.

- Use **admin redundancy synchronize boot-env** to synchronize the boot environments between the active and standby CPMs.

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 CPMs.

- Use **admin application-assurance upgrade** to load the new isa-aa image on the CPMs.
- Use **show application-assurance version** to verify new isa-aa image version running on the CPMs.
- 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 (**configure card** *slot-number* **mda** *mda-slot* **shutdown**) before an active ISA card is reset.

- Use **configure card** *slot-number* **mda** *mda-slot* **shutdown** to shut down an ISA card
- Use **clear mda** *slot-number/mda-slot* to reset an ISA card
- Use **configure card** *slot-number* **mda** *mda-slot* **no shutdown** to enable an ISA card
- Use **show application-assurance version** to verify the isa-aa.tim signatures version loaded on the CPMs 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 CPMs 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 ISSU Upgrade Procedure

This section describes the ISSU Upgrade Procedure which can be used:

- on routers with redundant CPMs.
- for Major ISSU across two major releases on 7450 ESS-7/12, 7750 SR-7/12, 7750 SR-12e, 7750 SR-a4/a8, 7750 SR-1e/2e/3e, 7950 XRS-16c/20/20e/40
- for Major ISSU across a single major release on 7450 ESS-7/12, 7750 SR-7/12, 7750 SR-12e, 7750 SR-a4/a8, 7750 SR-1e/2e/3e, 7950 XRS-16c/20/20e/40
- for Major ISSU across two major releases on routers running Release 15.0.R4 to 15.0.R16
- for Major ISSU across one major release on routers running Release 16.0.R4 to 16.0.R11
- for Minor ISSU on all platform types except 7750 SR-1 and SR-1s
- for Minor ISSU on routers running Release 19.10.R1 or higher

As of Release 19.10.R1, Minor ISSU is now supported when c**onfigure system management-interface configuration-mode** is set to **model-driven** or **mixed**. In the following procedures the **bof** and **file** commands need to be entered using the classic CLI.

Major ISSU is not supported when **configure system management-interface configuration-mode** is set to **model-driven** or **mixed**.

ISSU upgrade Procedure cannot be used if any manual firmware upgrade is required (such as **admin reboot upgrade**).

If any of the above criteria do not apply, the Standard Software Upgrade Procedure must be performed. See ISSU sub-section of Known Limitations for details.

ISSU limitations listed under Known Limitations should be taken into account for planning purposes before the ISSU is 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 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. It is also recommended to connect to the CONSOLE port on both CPMs prior to starting the ISSU.

The ISSU procedure is split into the following two phases:

- Phase A: Preparation and CPM Upgrade, with one procedure common to Minor and Major ISSU
- Phase B: Completion of the ISSU, with different procedures for Minor and Major ISSU

9.3.1 Phase A: Preparation and CPM Upgrade

Phase A covers ISSU preparation and the update of the CPM software.

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 the classic CLI **configure li save** command if the router is in classic or mixed configuration mode, or the MD-CLI command **admin save li** if the router is in model-driven configuration mode, and then back up the li.cfg file.

Step 2. Copy SR OS Images to cf3:

The SR OS image files must be copied to the cf3: device of the active CPM (only on the master chassis for 7950 XRS-40). 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:\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.

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 (only on the master chassis for 7950 XRS-40).

Step 4. Modify the Boot Options File to Point to 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. Synchronize Boot Environment

Once the Boot Options File has been modified, the active and standby CPM boot environments must be synchronized. As of Release 15.0.R1, this also synchronizes the boot environments on CPM C and CPM D of the extension chassis of a 7950 XRS-40.

- Use **admin redundancy synchronize boot-env** to synchronize the boot environments between the active and standby CPMs.

Step 6. Reboot the Standby CPM

In the sample output below, the active CPM is in Slot A and the standby CPM is in Slot B. Before performing ISSU on systems with CPMs, the **show card** output will display the information similar to the following:

```
A:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    up/active
B         cpm5                                     up    up/standby
===============================================================================
```

7950 XRS-40 systems will also display the extension CPMs on the extension chassis:

```
A:router1# show card
===========================================================================
Card Summary
===========================================================================
         Slot   Provisioned Type                        Admin Operational   Comments
                    Equipped Type (if different)        State State
         ---------------------------------------------------------------------------
         ...
         C      cpm-x20                                 up    up/ext-actv
         D      cpm-x20                                 up    up/ext-stby
         ...
```

Use **admin reboot standby now** to reboot the standby CPM and start the ISSU process.

After rebooting the standby CPM, the **show card** output will display information similar to the following:

```
A:router1# admin reboot standby now
A:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    up/active
B         cpm5                                     up    down/standby
===============================================================================
```

The extension CPMs on 7950 XRS-40 systems will initially stay in the up state:

```
A:router1# show card
===========================================================================
Card Summary
===========================================================================
Slot   Provisioned Type                        Admin Operational   Comments
           Equipped Type (if different)        State State
---------------------------------------------------------------------------
...
C      cpm-x20                                 up    up/ext-actv
D      cpm-x20                                 up    up/ext-stby
...
```

Step 7. Wait for Standby CPM to Synchronize

After the ISSU has been initiated, the card status of the standby CPM (in Slot B in this example) will show as “synching”. The standby CPM may only be in this synchronizing state for a brief period (depending on the amount of data that needs to be synchronized).

```
A:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    up/active
B         cpm5                                     up    synching/standby
===============================================================================
```

When the standby CPM has completely synchronized, the standby CPM will indicate a state of “ISSU”.

```
A:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    up/active
B         cpm5                                     up    ISSU/standby
==============================================================================
```

At this point, on systems with c**onfigure system management-interface configuration-mode** set to **model-driven**, it is not permitted to change the setting of **configure system management-interface configuration-mode**.

At this point, on 7950 XRS-40 systems, SR OS will automatically attempt to reboot the extension standby CPM to bring it up with the new software image. The automatic reboot of the extension standby CPM is triggered whenever a master standby CPM comes online and the system sees that the card has entered or exited an ISSU state. Log events will indicate that the extension standby CPM is attempting to be rebooted or if the system is unable to attempt the reboot (for example, loss of connectivity to the extension chassis). The extension standby CPM will reboot and initially show a state of “provisioned”:

```
A:router1# show card
==========================================================================
Card Summary
==========================================================================
Slot   Provisioned Type                        Admin Operational   Comments
           Equipped Type (if different)        State State
---------------------------------------------------------------------------
.....
A      cpm-x20                                 up    up/active
B      cpm-x20                                 up    ISSU/standby
C      cpm-x20                                 up    up/ext-actv
D      cpm-x20                                 up    provisioned/*
           (not equipped)
===========================================================================
* indicates that the corresponding row element may have been truncated.
```

After a few moments, the extension standby CPM will transition to an ISSU state:

```
A:router1# show card
==========================================================================
Card Summary
==========================================================================
Slot   Provisioned Type                        Admin Operational   Comments
           Equipped Type (if different)        State State
---------------------------------------------------------------------------
.....
A      cpm-x20                                 up    up/active
B      cpm-x20                                 up    ISSU/standby
C      cpm-x20                                 up    up/ext-actv
D      cpm-x20                                 up    ISSU/ext-stby
===========================================================================
```

If the extension standby CPM does not transition into the ISSU state then a manual **clear card** *m* (where *m* is the card letter of the current extension standby CPM) could be attempted or the ISSU could be aborted. In this case, point the BOF back to the old images, put the old boot.ldr back in the root directory of CF3: on the master active CPM, synchronize the boot environment, reboot (**clear card**) the extension standby CPM, and finally reboot the master standby CPM (**admin reboot standby**).

Warning: On 7950 XRS-40 systems, an extension CPM that is in an ISSU state cannot become an extension active CPM. The extension chassis temporarily has no CPM redundancy while in this state, so the operator should move on to the next step as soon as possible.

If the extension active CPM reboots (or goes down for any reason) while the extension standby CPM is in an ISSU state, then all communications with the extension chassis will be lost (the extension standby CPM in the ISSU state cannot take over as the extension active CPM), resulting in a loss of service for the extension shelf and reduced fabric capacity for the master shelf.

If the extension active CPM boots back up, it will come up in the new software image as an extension standby (it cannot be an extension active CPM due to software mismatch with the master active CPM).

See Step 8 for more information.

Step 8. If necessary, to recover from the above situation on the 7950 XRS-40, the operator should perform the following steps.

i. Point the BOF back to the old images.

ii. Put the old boot.ldr back in the root directory of CF3: on the master active CPM.

iii. Synchronize the boot environment.

iv. Reboot the master standby CPM (**admin reboot standby**). This will also cause the system to reboot the extension CPMs, resulting in all four CPMs being up and in the original software image.

9.3.2 Phase B: Completion of the ISSU

Phase B of the ISSU procedure is different for Minor ISSU and Major ISSU. Proceed to the appropriate procedure.

- Phase B (Minor ISSU)
- Phase B (Major ISSU)

9.3.2.1 Phase B (Minor ISSU)

The following steps describe Phase B of the ISSU procedure for Minor ISSU.

Step 1. (Minor ISSU) Switchover the CPM

After the standby CPM has synchronized (Operational State = ISSU/ standby), then the operator can proceed to the next phase of Minor ISSU.

Note that if the standby CPM is being held in the “down” operational state, then log 99 can be checked for log events that will explain the reason. For example, the following two log events would be in the log if the system contains deprecated hardware (such as the m4-choc3-sfp):

```
122 2015/05/30 16:21:03.83 EDT MAJOR: CHASSIS #2001 Base
Card B "Class CPM Module : failed, reason: Issu
Unsupported Scenario, No Reload"
```

```
121 2015/05/30 16:21:03.84 EDT MAJOR:CHASSIS #2001 Base
Card B "Class CPM Module : failed, reason: Unsupported
MDA type m4-choc3-sfp in slot 1/2"
```

After the standby CPM has synchronized and indicates a card status of “ISSU”, a CPM switchover (from A to B in this example) must be performed in order to force the CPM running the new software image to become the active CPM. The switchover command will cause the active CPM to reboot.

- Use **admin redundancy force-switchover** to make the CPM with the new software image become the active CPM.

Note that, if c**onfigure system management-interface configuration-mode** is set to **model-driven**, and the active CPM reboots for any reason other than the **force-switchover** command, then the ISSU will be terminated and a full node reboot will occur.

Note that, when the switchover command is issued, a warning will be printed if any cards are equipped:

```
WARNING: After switchover, the following resets will be needed:
```

For each IOM/IMM that is equipped, regardless of state, a one-line summary is displayed to indicate whether the card will be hard reset or Soft Reset, along with a reason for the hard reset. See Step 1 of the Major ISSU procedure for more details.

For systems with **configure system management-interface configuration-mode** set to **model-driven**, the configuration is transferred from the active CPM to the standby before the active CPM releases control. Depending on the size of the configuration, this can last multiple minutes. After the switchover, the newly active CPM performs an integrity and validation check on the configuration. Any errors will result in a reboot of the newly active CPM.

Step 2. (Minor ISSU) If Necessary, Re-establish a Console Session

If the ISSU is performed from the serial port CONSOLE on the CPM and there is only one terminal available (that is, one PC with a serial port), the console session must be re-established on the newly active CPM.

Step 3. (Minor ISSU) Wait for Standby CPM to Synchronize

Before continuing with the ISSU procedure, the standby CPM must re-synchronize by transitioning from “down”, to “synchronizing”, and finally to the “up” state. Use the **show card** command to monitor the status of the IOMs and IMMs. Note that the IOMs and IMMs now have an “ISSU” status indicating that the active CPM is running the new image, as in this example for systems equipped with CPMs.

```
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    ISSU
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    ISSU
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    down/standby
B         cpm5                                     up    up/active
===============================================================================
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    ISSU
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    ISSU
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    synching/standby
B         cpm5                                     up    up/active
===============================================================================
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    ISSU
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    ISSU
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
```

```
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

Step 4. (Minor ISSU) Reset the line cards to Load the New Image

The IOMs, ISMs, and IMMs must now be reset to load the new image. If the cards will be Soft Reset (see below), see ISSU in the Known Limitations section for the source/starting release of the upgrade. Soft Reset limitations should be taken into account for planning purposes before the ISSU is performed.

- Use the **clear card** *n* **soft hard-reset-unsupported-mdas** command to Soft Reset a line card. The line card data path and MDAs/ISAs are not reset in Soft Reset compatible cases, resulting in a very brief service interruption.
- If the Soft Reset is blocked, then use the **clear card** *n* command to hard reset the line card. This will reboot the line card and its MDAs and ISAs, causing an outage for the duration of the reboot.

Note:

- The system does not allow cards to run in an ISSU state indefinitely; the system automatically hard resets the IOMs/IMMs/ISMs after two hours. The “Comments” field in the show card state output displays the time until the system resets the line cards in the ISSU state.
- It is recommended to Soft Reset no more than one line card at a time to ensure that the line card download process does not impact control plane protocols. Wait for the operational state to be “up” before proceeding to the next line card.
- With the Deferred MDA Reset enhancement (introduced in Release 10.0.R1), Soft Reset of a card is allowed to proceed even when the MDA firmware does not match the MDA firmware in the new image. The operator is informed of MDAs running below the latest revision of firmware with CHASSIS log event #2082. The MDA can be upgraded to the latest firmware (after the Soft Reset) by performing a hard reset of the MDA (**clear mda** *x/y*).

The sample output below shows the operational state transition for a single line card.

```
B:SoftReset1# clear card 4 soft hard-reset-unsupported-mdas
A: SoftReset1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
```

```
1         iom5e:he1200g+                           up    ISSU
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    soft reset
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

When the IOM/IMM/ISM is in the “up” state, it will have the new image so it will no longer have an “ISSU” operational state as shown in the sample output below.

```
B: SoftReset1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    ISSU
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

The timing and order of the line card resets should be sequenced to maximize the effectiveness of any redundant interfaces (LAGs, VRRP, etc.) spanning IOM/IMM/ISM, MDA, or any ISA redundancy deployed slots.

The sample output below shows the operational state transitions for a single IOM.

```
B:router1# clear card 1
B: router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    provisioned
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
```

```
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

When the line card is in the “up” state, it will have the new image so it will no longer have an “ISSU” operational state as shown in the sample output below.

```
B: router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    ISSU
3         imm4-100gb-cfp4                          up    ISSU
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    ISSU
6         imm-2pac-fp3                             up    ISSU
7         imm-2pac-fp3                             up    ISSU
8         iom3-xp-c                                up    ISSU
9         iom5-e:he1200g+                          up    ISSU
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

When all of the line cards have been rebooted, the ISSU is complete. It is recommended to save the configuration (admin save) after an upgrade has been performed and the system is operating as expected. This will ensure that all configurations are saved in a format that is fully compatible with the newly running release.

9.3.2.2 Phase B (Major ISSU)

The following steps describe Phase B of the ISSU procedure for Major ISSU.

Step 1. (Major ISSU) Switch Over the CPM

After the standby CPM has synchronized (Operational State = ISSU/ standby), then the operator can proceed to the next phase of Major ISSU.

Note that if the standby CPM is being held in the “down” operational state, look at log 99 for log events that explain the reason. For example, if the system contains deprecated hardware such as the m4-choc3-sfp:

```
122 2015/05/30 16:21:03.83 EDT MAJOR:
CHASSIS #2001 Base Card B "Class CPM Module :
failed, reason: Issu Unsupported Scenario, No Reload"
121 2015/05/30 16:21:03.84 EDT MAJOR:
CHASSIS #2001 Base Card B "Class CPM Module :
failed, reason: Unsupported MDA type m4-choc3-
sfp in slot 1/2"
```

After the standby CPM has synchronized and indicates a card status of “ISSU/standby”, a CPM switchover (from A to B in this example) must be performed in order to force the CPM running the new software image to become the active CPM. The **switchover** command will cause the active CPM to reboot.

On 7950 XRS-40 systems, the **force-switchover** command also causes an extension CPM switchover in the extension chassis (and the previous extension active CPM reboots). The **force-switchover** is not allowed unless:

- the status of all CPM IcPorts are up and
- both the master standby CPM and extension standby CPM are in the ISSU state

Use **admin redundancy force-switchover** to make the CPM with the new software image become the active CPM.

Note that if the active CPM reboots for any reason other than the **force-switchover** command, then the ISSU will be terminated and a full node reboot will occur.

When the switchover command is issued, a warning will be printed if any cards are equipped:

```
WARNING: After switchover the following HARD and SOFT resets will occur:
```

For each line card that is equipped, regardless of its state, a one-line summary is displayed to indicate whether the card will be hard reset or Soft Reset, along with a reason for the hard reset. The following example shows a particular card and MDA configuration, along with the resulting ISSU hard reset or Soft Reset reasons.

```
A:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
```

```
6         imm-2pac-fp3                              up    up
7         imm-2pac-fp3                              up    up
8         iom3-xp-c                                 up    up
9         iom5-e:he1200g+                           up    up
A         cpm5                                      up    up/active
B         cpm5                                      up    ISSU/standby
===============================================================================

A:router1# show mda
===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned Type                            Admin     Operational
                Equipped Type (if different)            State     State
-------------------------------------------------------------------------------
1     1     me12-100gb-qsfp28                           up        up
      2     me12-100gb-qsfp28                           up        up
2     1     me12-100gb-qsfp28                           up        up
      2     me12-100gb-qsfp28                           up        up
3     1     m4-100g-cfp4                                up        up
4     1     m4-100g-cfp4                                up        up
5     1     m4-100g-cfp4                                up        up
6     1     p1-100g-cfp                                 up        up
      2     p1-100g-cfp                                 up        up
7     1     p1-100g-cfp                                 up        up
      2     p1-100g-cfp                                 up        up
8     1     m4-10gb-xp-xfp                              up        up
      2     m4-10gb-xp-xfp                              up        up
9     1     me6-100gb-qsfp28                            up        up
      2     me12-100gb-qsfp28                           up        up

A:router1# admin redundancy force-switchover
WARNING: After switchover the following HARD and SOFT resets will occur:
IOM 1: SOFT (MDAs: 1/1 SOFT, 1/2 SOFT)
IOM 2: SOFT (MDAs: 2/1 SOFT, 2/2 SOFT)
IOM 3: SOFT (MDAs: 3/1 SOFT)
IOM 4: SOFT (MDAs: 4/1 SOFT)
IOM 5: SOFT (MDAs: 5/1 SOFT)
IOM 6: SOFT (MDAs: 6/1 SOFT, 6/2 SOFT)
IOM 7: SOFT (MDAs: 7/1 SOFT, 7/2 SOFT)
IOM 8: SOFT (MDAs: 8/1 SOFT, 8/2 SOFT)
IOM 9: SOFT (MDAs: 9/1 SOFT, 9/2 SOFT)
```

The reason codes are as follows:

- unsupported: Soft Reset not supported on the assembly
- incompatible: the specific upgrade scenario being attempted (from software image X to software image Y) is not Soft Reset compatible (for example: mandatory datapath firmware upgrades on an MDA or IMM)
- offline: the assembly is not currently operational
- not present: the card is not present
- any MDA/XMA hard reset forces IOM/XCM hard reset: one of the MDAs/XMAs cannot be upgraded without IOM/XCM hard reset

No reason codes are given for MDAs/XMAs that are shut down (a reset of those MDAs/XMAs will have no impact on service), or for the second MDA identifier in a slot that contains an IMM.

After the IOM/XCM summary, the following prompt is given to the operator:

```
WARNING: Major in service software upgrade in progress.
Are you sure you want to switchover (y/n)?
```

The switchover may be blocked in various error scenarios. A warning will explain the problem. For example, the following message will occur if the standby does not have enough compact flash space for the configuration to be synchronized:

```
MINOR: CHMGR #1055 - Major ISSU sync of config to standby failed
```

If the switchover is attempted when the standby is not in an “ISSU/standby” state, then normal High-Availability switchover behavior will apply.

Step 2. (Major ISSU) If Necessary, Re-establish a Console Session

If the ISSU is performed from the serial port CONSOLE on the CPM, and there is only one terminal available (i.e., one PC with a serial port), the console session must be re-established on the newly active CPM.

Step 3. (Major ISSU) Line Card Update

When the switchover command is used in Major ISSU, the active CPM will prepare the system for the ISSU and then reboot. The other CPM (previously the standby and running the newer software load) will take over as the active CPM.

After the switchover, a command prompt will be available on the newly active CPM. Configuration changes are not allowed at this point, but most **show**, **clear** and **admin** commands are available. If the operator attempts to use a command that is invalid during this phase, they will receive the following error:

```
*B:Dut-A# configure service epipe 3 customer 1 create
MINOR: CLI Command not allowed while becoming active.
```

After the Major ISSU is complete, the full CLI functionality will be available.

Shortly after the switchover, all line cards are reset so that they can upgrade to the new image. The reset will be a Soft Reset for any supported combinations of cards, and hard reset for all other cases (with reasons displayed for each line card as described in previous steps).

Note:

- The Soft Reset section of the Known Limitations for the source/starting release of the upgrade should be taken into account for planning purposes before the ISSU is performed.

The sample output below shows the operational state transition for the cards in the system after the CPM running the new software image first takes over:

```
*B:router1# show redundancy synchronization
===========================================================================
Synchronization Information
===========================================================================
Standby Status               : disabled
Last Standby Failure         : N/A
Standby Up Time              : N/A
Standby Version              : N/A
Failover Time                : 05/30/2015 16:00:33
Failover Reason              : user forced switchover
Boot/Config Sync Mode        : None
Boot/Config Sync Status      : No synchronization
Last Config File Sync Time   : Never
Last Boot Env Sync Time      : Never
Rollback Sync Mode           : None
Rollback Sync Status         : No Rollback synchronization
Last Rollback Sync Time      : Never
===========================================================================
```

```
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    soft reset
              (not equipped)
2         iom5e:he1200g+                           up    soft reset
              (not equipped)
3         imm4-100gb-cfp4                          up    soft reset
              (not equipped)
4         imm4-100gb-cfp4                          up    soft reset
              (not equipped)
5         imm4-100gb-cfp4                          up    soft reset
              (not equipped)
6         imm-2pac-fp3                             up    soft reset
              (not equipped)
7         imm-2pac-fp3                             up    soft reset
              (not equipped)
8         iom3-xp-c                                up    soft reset p
              (not equipped)
9         iom5-e:he1200g+                          up    soft reset
              (not equipped)
A         cpm5                                     up    down/standby
              (not equipped)
B         cpm5                                     up    up/active
===============================================================================
```

The new extension standby CPM on 7950 XRS-40 systems will initially be in the “provisioned” state:

```
===========================================================================
Card Summary
```

```
===========================================================================
Slot   Provisioned Type                     Admin    Operational   Comments
           Equipped Type (if different)     State    State
---------------------------------------------------------------------------
A      cpm-x20                              up       down/standby
           (not equipped)
B      cpm-x20                              up       up/active
C      cpm-x20                              up       provisioned/*
           (not equipped)
D      cpm-x20                              up       up/ext-actv
```

A few seconds later, most of the cards have been detected and are in the Soft Reset or booting state. The standby CPM will remain as “down/ standby” until all Soft Resets are completed.

```
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    soft reset
2         iom5e:he1200g+                           up    booting
3         imm4-100gb-cfp4                          up    soft reset
4         imm4-100gb-cfp4                          up    soft reset
5         imm4-100gb-cfp4                          up    booting
6         imm-2pac-fp3                             up    soft reset
7         imm-2pac-fp3                             up    soft reset
8         iom3-xp-c                                up    soft reset p
9         iom5-e:he1200g+                          up    booting
A         cpm5                                     up    down/standby
B         cpm5                                     up    up/active
===============================================================================
```

The following output shows the cards having completed their resets and are now running with the new software image. The standby CPM will synchronize with the active CPM once all Soft Resets are completed.

```
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    synching/standby
B         cpm5                                     up    up/active
===============================================================================
```

Step 4. (Major ISSU) ISSU Completion

Monitor the node to ensure that it returns to normal operation. All line cards should return to the “up” state, and the standby CPM should return to the operational “up” state. Note that the standby CPM may spend a few minutes in the synchronizing state before finally settling in the “up” state.

The following output shows the IOM/IMMs backed up, and the standby CPM synchronized (“up”).

```
B:router1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned Type                         Admin Operational   Comments
              Equipped Type (if different)         State State
-------------------------------------------------------------------------------
1         iom5e:he1200g+                           up    up
2         iom5e:he1200g+                           up    up
3         imm4-100gb-cfp4                          up    up
4         imm4-100gb-cfp4                          up    up
5         imm4-100gb-cfp4                          up    up
6         imm-2pac-fp3                             up    up
7         imm-2pac-fp3                             up    up
8         iom3-xp-c                                up    up
9         iom5-e:he1200g+                          up    up
A         cpm5                                     up    up/standby
B         cpm5                                     up    up/active
===============================================================================
```

The new extension standby CPM on 7950 XRS-40 systems will finally transition to an up state:

```
===========================================================================
Card Summary
===========================================================================
Slot   Provisioned Type                     Admin    Operational   Comments
           Equipped Type (if different)     State    State
---------------------------------------------------------------------------
A      cpm-x20                              up       up/standby
B      cpm-x20                              up       up/active
C      cpm-x20                              up       up/ext-stby
D      cpm-x20                              up       up/ext-actv


*B:Dut-A# show redundancy synchronization
===========================================================================
Synchronization Information
===========================================================================
Standby Status               : standby ready
Last Standby Failure         : N/A
Standby Up Time              : 2015/05/30 16:05:03
Standby Version              : ...<version info>...
Failover Time                : 05/30/2015 16:00:33
Failover Reason              : user forced switchover
Boot/Config Sync Mode        : None
Boot/Config Sync Status      : No synchronization
Last Config File Sync Time   : Never
Last Boot Env Sync Time      : Never
Rollback Sync Mode           : None
Rollback Sync Status         : No Rollback synchronization
Last Rollback Sync Time      : Never
===========================================================================
```

When all of the line cards have been rebooted, and the active and standby CPMs are synchronized, the ISSU is complete. Full CLI functionality will be available at this point.

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

Step 5. (Major ISSU) Optional Post-ISSU Actions

With the Deferred MDA Reset enhancement, Soft Reset of a card is allowed to proceed even when the MDA/XMA firmware does not match the MDA/XMA firmware in the new image. The operator is informed of MDAs/XMAs running below the latest revision of firmware with CHASSIS log event #2082. The MDA/XMA can be upgraded to the latest firmware (after the Soft Reset) by performing a hard reset of the MDA/XMA (**clear mda** *x*/*y*).

9.4 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
- 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 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 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. For 7450 ESS, 7750 SR, and 7950 XRS systems that have installed license files, it is recommended to update the license file to cover both the active release and the target release in advance of any upgrade. Use CLM to download and activate a new license file to cover the two releases. Confirm that the target release is available in the node using the **show**>**system**>**license available-licenses** command.

Step 2. 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.

When saving the configuration at the start of a software upgrade, the **detail** option of **admin save** (for example, **admin save detail**) should not be used on the VSR platform. Upgrades with **admin save detail** are not supported on those platforms.

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 3. Copy the SR OS images to cf3:

The SR OS image files must to be copied to the cf3: device on the active CPM (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.

Step 4. 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 (only on the master chassis for 7950 XRS-40).

Step 5. 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 6. For Redundant CPMs, synchronize boot environment

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

As of Release 15.0.R1, this command also synchronizes the boot environments on CPM C and CPM D of the extension chassis of a 7950 XRS-40. If using a release prior to Release 15.0, then the boot environment of CPM C and CPM D must be synchronized manually. See Upgrading to Release 19.5.R1 or Higher for these steps.

Step 7. Reboot the chassis

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

Step 8. 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 19.10.R6 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 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 IOM, IMM, ISM, XCM, MDA, and XMA modules. Even when only one line card is deployed, impedance panels are required.

10.4 Multiservice Integrated Services Adapter (ISA)

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

Table 20 **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)	✓	✓	✓	✓	✓	✓	✓
Retransmission, Fast Channel Change, and Video Quality Monitoring (isa-video/isa2-video)	✓	✓	✓	✓			
Tunnel Services, including IPsec (isa-tunnel/isa2-tunnel)	✓[1]	✓		✓		✓	
Network Address Translation (isa-bb/isa2-bb)	✓	✓	✓	✓	✓	✓	✓
L2TP LNS Service (isa-bb/isa2-bb)	✓	✓	✓	✓	✓	✓	✓
WLAN-GW (isa-bb/isa2-bb)	✓	✓	✓	✓	✓	✓	✓

Note:

1. MS-ISA only. Not supported on MS-ISA-E.

Table 21 **Compatible 7450 ESS 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**
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.

10.5 Compact Flash Devices

- Only Nokia-sourced compact flash devices for the SR OS are supported.
- Nokia recommends that the compact flash in the CF3 slot be at least 2 GB. 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.

10.6 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.

10.7 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 requires a mandatory reboot of the chassis.

- Some commands use the forwarding plane (FP) number as an identifier for the FP on the card. On FP4 XMAs, there can be more than two FPs assigned to a specific card slot. For the 7950 XRS-20/20e XCM2, the XMA in MDA slot 1 uses FP numbers 1 and 2 and the XMA in MDA slot 2 uses FP numbers 3 and 4. For the 7750 SR-14s XCM-s, the XMA-s in MDA slot 1 uses FP numbers 1 to 4 and the XMA-s in MDA slot 2 uses FP numbers 5 to 8.

10.8 PXC

- PXC is supported only on Ethernet ports.
- PXC is not supported on ports in DWDM, WAN or OTN mode.
- PXC is not supported on iom4-e-hs (HSQ) and HS-MDAv2.
- PXC is supported on break-out ports (connector-based ports) with the exception of the 10 x 10GE breakout (that is, **c10-10g**)
- PXC is supported on 10G and above with exception of 40G ports on FP3-based hardware. Table 22 shows PXC support by platform.

Table 22 **PXC Support by Platform**

Platform	Support Description	Exceptions	CLI String of Exception
7750 SR-a4/a8	10G ports and above	Exception is the ports on the MDA-a 7750 SR 4pt 10GE SFP+	ma4-10gb-sfp+
7750 SR-1e/2e/3e	10G ports and above	Exception is the ports at 40G speed on the me2-100gb-qsfp28 module	me2-100gb-qsfp28

Table 22 **PXC Support by Platform (Continued)**

Platform	Support Description	Exceptions	CLI String of Exception
7750 SR-7/12/12e and SR line cards in ESS chassis configured as "7450 Mixed-Mode"	Supported on CPM5, chassis mode D, and 10G ports and above	Exception is the ports at 40G speed on the me2-100gb-qsfp28 module p3-40G-QSFP on IMM-2PAC-FP3 IMM	me2-100gb-qsfp28 p3-40g-qsfp
7750 SR-1/1s/2s/7s/14s	10G and above Note that these platforms are FP4 only	--	--

10.9 Flex PW-port

- When dual-homing with SRRP is used with Flex PW-port, the SRRP messaging path between the two BNGs must be symmetric. The downstream and upstream (returning) SRRP messages must traverse the same path between the two BNGs. An asymmetric upstream/downstream path created (for example, via strict hops) may lead during link failures to undesired scenarios where both BNGs transition into the master state.

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. Obsoleted configuration executed via an older configuration file will be skipped with a message provided.
- 7210 SAS Ethernet satellites configured for IEEE 1588 Transparent Clock functionality are intended to be used with direct fiber connectivity having fixed and symmetric delay between the SR OS host and the satellite. Use of intervening networks, including optical transport networks, can impact performance and should be verified before deployment.

10.11 Multi-Chassis Synchronization (MCS)

- MCS for subscriber management applications is supported between following chassis types:
  - ESS-7/12 and ESS-7/12
  - SR-a4/a8 and SR-a4/a8
  - SR-1e/2e/3e and SR-1e/2e/3e
  - SR-1 and SR-1
  - SR-7/12/12e and SR-7/12/12e
  - SR-1s and SR-1s
  - SR-2s and SR-2s
  - SR-7s and SR-7s
- 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.
- MCS and SRRP between a redundant BNG pair is supported for nodes running different software versions of up to two major SR OS releases. While the nodes of the redundant BNG pair are running different major software versions, only the feature set of the lowest software version can be used. The period in which different SR OS releases are deployed should be kept to a minimum. Nokia recommends using major ISSU to upgrade the earlier SR OS version.

  Supported MCS applications for an upgrade between two major SR OS releases:
  - **diameter-proxy**
  - **igmp**
  - **l2tp**
  - **local-dhcp-server**
  - **mld**
  - **port**
  - **python**
  - **srrp**
  - **sub-host-trk**
  - **sub-mgmt**
  - **igmp-snooping**

10.12 Telemetry/gRPC

- The SR OS gRPC Telemetry interface in Release 19.7.R1 is based on OpenConfig gnmi.proto version 0.7.0.

10.13 ATM

- 7750 SR and 7450 ESS 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]

10.14 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 do not advertise the protocol field compression (PFC) option during LCP negotiation. [25996, 29923]

10.15 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.

10.16 TCP Authentication Extension

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

10.17 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.

10.18 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.

10.19 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]

10.20 IS-IS TE

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

10.21 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.

10.22 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]

10.23 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.

10.24 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]

10.25 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.

10.26 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]

10.27 PIM

- To ensure proper GRT/VRF extranet functionality, it is strongly recommended 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 recommended 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]

10.28 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.

10.29 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, 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 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.

10.30 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 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 16.0.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 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 Removed Commands in Model-driven Interfaces for more information.
- In Release 16.0.R1, the *name* parameter (*admin-name* for ETH-CFM) on the classic CLI **create** line of the following objects is immutable. The *name* (and *admin-name*) cannot be changed after creating the object (without deleting and recreating the object):
  - all services (**configure service vprn**, **vpls**, **epipe**, etc.)
  - **configure mirror mirror-dest**
  - **configure service pw-template** contexts
  - **configure service customer**
  - **configure filter ip-filter** | **ipv6-filter** | **mac-filter**
  - **configure qos network** |**sap-ingress** | **sap-egress**
  - **configure eth-cfm domain** | **association**

  During an upgrade to Releases 16.0.R1 or higher, any object in this list that does not have a *name* configured prior to the upgrade, will be assigned a *name* created from converting their numerical ID (**service-id**, **filter-id**, etc.) to a string. For example:

```
configure filter ip-filter 35 name "35"
```

The *name* parameter is used in model driven (MD) interfaces (for example, the MD-CLI, NETCONF, gRPC) as the key of the object (instead of using the numerical ID as the key). For example:

```
configure filter ip-filter "local-sites-23a"
```

See Software Upgrade Procedures for more information.

10.31 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.

10.32 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 **reliable** 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]

10.33 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]
- The EVI-derived route-target (**bgp-evpn**>**evi**) is only valid in a service BGP instance if **bgp-evpn** is enabled in that instance. In a service where EVPN and BGP-MH are both enabled in an instance with no explicit or BGP-AD derived route-targets, **bgp-evpn** should be enabled first. Otherwise, EVPN routes will be advertised with the default route-target (0:0) rather than the EVI-derived route-target. [299549]

10.34 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]

10.35 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]

10.36 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.

10.37 IPsec Compatibility

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

Table 23 **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 24 **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

10.38 Mirror Service

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

10.39 NAT

- Although L2-aware subscribers in ESM are synchronized using MCS in a multi-chassis environment, NAT pools and bindings are not. Furthermore, L2-aware NAT pools for the same ESM subscriber on two redundant nodes must contain different outside IP addresses. This results in change of the outside (NAT) IP address for a subscriber after a node or SRRP switchover. The user TCP/UDP sessions established before the switchover become stale and time out on the application level and the pool level (in case of the SRRP switchover). However, DHCP/PPPoE sessions on the client side are preserved.

10.40 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.

10.41 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, 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.

10.42 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.

10.43 BFD on LSPs

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

10.44 BFD VCCV

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

Table 25 **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

10.45 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 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 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.

10.46 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.

10.47 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 VSR Platform and Feature Support

This section describes topics specific to Virtualized Service Router (VSR).

- Supported NFV Infrastructure
- VSR Deployment Considerations
- Supported VSR Features
- SR OS Features not Supported on VSR

This section summarizes the key features supported in the Release 19.10.R6 of SR OS software that are applicable to the following key VSR applications:

- Provider Edge (PE)
- Data Center Gateway (DCGW)
- Broadband Network Gateway (BNG)
- L2TP Network Server (LNS)
- Security Gateway (SeGW)
- Application Assurance (AA)
- Network Address Translation (NAT)
- MAP-T Border Relay (MAP-T BR)
- WLAN Gateway (WLAN-GW)
- Virtual Residential Gateway (vRGW)
- BGP Route Reflector (RR)

11.1 Supported NFV Infrastructure

11.1.1 Server Hardware: CPUs

Nokia recommends Nokia Airframe servers for deployment of VSR virtual machines. However, VSR is supported on any server with one or more of the following CPUs:

- Intel Xeon E5-26*xx*-v2 CPUs (Intel Ivy Bridge)
- Intel Xeon E5-26*xx*-v3 CPUs (Intel Haswell)
- Intel Xeon E5-26*xx*-v4 CPUs (Intel Broadwell)
- Intel Xeon 5*xxx*/6*xxx*/8*xxx* Gold or Platinum CPUs (Intel Skylake-SP)

CPUs with at least 12 cores per socket/NUMA node are recommended.

VSR is not supported on any server powered by AMD or ARM CPUs.

11.1.2 Server Hardware: Memory

The server should be equipped with enough DRAM memory to meet the memory requirements of the host and enough memory to back the memory of each guest virtual machine without oversubscription. See VM Memory Requirements by Function Mix for more information.

11.1.3 Server Hardware: NICs

If virtual interfaces of the VSR virtual machine will use SR-IOV or PCI pass-through, the associated physical NICs in the server must be one of the types listed in Table 26 (for SR-IOV) or Table 27 (for PCI pass-through). Note that the list of supported NICs depends on the hypervisor used by the host.

Table 26 shows the supported NICs for SR-IOV. For SR-IOV, the maximum number of VFs per VSR is 20.

Table 26 Supported NICs for SR-IOV

Vendor	Model	Supported Speeds (Gb/s)	VSR on KVM Host	VSR on ESXi Host
Intel	X520-DA2	10	Yes	Yes Recommended driver: 3.7.13.7.14iov-20vmw.600.0.0. 2494585
Intel	X710-DA4, X710-DA2	10	Yes Recommended i40e driver version: 2.1.14-k	No

Table 26 **Supported NICs for SR-IOV (Continued)**

Vendor	Model	Supported Speeds (Gb/s)	VSR on KVM Host	VSR on ESXi Host
Mellanox	Connect-X4 MCX416ACCAT	40 100	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 12.16.1021	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 12.22.1002
Mellanox	Connect-X4 MCX4121AACAT	25	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 14.17.1011	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 14.17.1011
Mellanox	Connect-X4 MCX414ABCAT	40	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 12.16.1021	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 12.16.1021
Mellanox	Connect-X5 MCX516A-CCAT	100	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 16.21.2010	Yes Minimum firmware: 16.21.2010
Mellanox	ConnectX-5 Ex MCX516A-CDAT	100	Yes Minimum firmware: 16.24.1000	Yes Minimum firmware: 16.24.1000
HPE	2-port, 10/25GbE 817753-B21 817749-B21	10 25	Yes Driver: nmlx5_core version: 4.17.13.8 Minimum firmware: 14.18.2030	Yes Driver: nmlx5_core version: 4.17.13.8 Minimum firmware: 14.18.2030
Intel	XXV710-DA2	10 25	Yes	No

Table 27 shows the supported NICs for PCI Pass-through.

Table 27 **Supported NICs for PCI Pass-through**

Vendor	Model	Supported Speeds (Gb/s)	VSR on KVM Host	VSR on ESXi Host
Intel	X520-DA2	10	Yes	Yes Recommended driver: 3.7.13.7.14iov-20vmw.600.0.0.2494585
Intel	X710-DA4, X710-DA2	10	Yes Recommended i40e driver version: 2.1.14-k Recommended firmware version: 6.01	No
Intel	X722	10	Recommended firmware version: 3.1d or 3.2d	No
Mellanox	Connect-X3 MCX313A-BCCT MCX353A-FCCT	40	Yes	No
Mellanox	Connect-X3 MCX314A-BCCT MCX-354A-FCCT	40	Yes, but only the first port is usable	No
Mellanox	Connect-X4 MCX416A-CCAT	40 100	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 12.16.1021	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 12.22.1002

Table 27 **Supported NICs for PCI Pass-through (Continued)**

Vendor	Model	Supported Speeds (Gb/s)	VSR on KVM Host	VSR on ESXi Host
Mellanox	Connect-X4 MCX4121A-ACAT	25	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 14.17.1011	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 14.17.1011
Mellanox	Connect-X4 MCX414ABCAT	40	Yes Minimum driver: 3.3-1.0.0.0 Minimum firmware: 12.16.1021	Yes Recommended driver: 4.16.10.3-1OEM.650.0.0.4598673 Recommended firmware: 12.16.1021
Mellanox	Connect-X5 MCX516A-CCAT	100	Yes Minimum firmware: 16.21.2010	Yes Minimum firmware: 16.21.2010
Mellanox	ConnectX-5 Ex MCX516A-CDAT	100	Yes Minimum firmware:16.24.1000	Yes Minimum firmware:16.24.1000
HPE	2-port, 10/25GbE 817753-B21 817749-B21	10 25	Yes Driver: nmlx5_core version: 4.17.13.8 Minimum firmware: 14.18.2030	Yes Driver: nmlx5_core version: 4.17.13.8 Minimum firmware: 14.18.2030
Intel	XXV710-DA2	10 25	Yes	No

11.1.4 CPU Pinning

CPU pinning is a hypervisor action that locks the execution of each guest virtual CPU (vCPU) to a physical CPU (pCPU) core/thread. CPU pinning is required for all VSR guests.

When hyperthreading is disabled in the BIOS of the host machine, each vCPU of a single VSR VM must be mapped to its own isolated pCPU core on the same NUMA node, but there is no required order to the pCPU cores. The pCPU cores do not have to be consecutive.

When hyperthreading is enabled in the BIOS of the host machine, each vCPU of a single VSR VM must be mapped to its own isolated pCPU thread, and the mapping order is important. The guest VM must have an even number of vCPUs. The first two vCPUs (0 and 1) must be pinned to two sibling threads of the same pCPU core. The next two vCPUs (2 and 3) must be pinned to two siblings on some other pCPU core of the same NUMA node, and so on.

Refer to the *VSR Setup and Installation Guide* for hypervisor-specific instructions on how to meet these CPU pinning requirements.

11.1.5 CPU Isolation

The host OS must not schedule its own tasks to the CPU cores assigned to a VSR VM. VSR performance and stability could be compromised.

On Linux KVM hosts, this can be achieved using one of two methods: using the **isolcpus** kernel boot parameter or the **systemd** CPU affinity. The **isolcpus** kernel parameter specifies a list of CPU cores that should be avoided by the host scheduler, and this should match the list of CPUs to which VSR VMs are pinned. The CPU affinity setting in /etc/systemd/system.conf specifies the subset of CPU cores that must be used by systemd for its tasks. The list should not overlap with the list of CPUs to which VSR VMs are pinned.

On VMware ESXi hosts, CPU isolation can be achieved using one of two methods: by setting the **latencySensitivity** property to **high**, or by editing the VMX file to include specific vCPU affinity directives.

Refer to the *VSR Setup and Installation Guide* for more information.

11.1.6 Server BIOS Settings

The following BIOS settings are mandatory on a compute server hosting any type of VSR VM (regardless of the VSR application or performance requirements).

- Intel VT-x must be enabled.
- NUMA (Non Uniform Memory Access) must be enabled.
- SR-IOV must be enabled if this technology is planned for use. The actual BIOS setting may be named “SR-IOV Global Enable”.
- Intel VT-d must be enabled if either SR-IOV or PCI pass-through is planned for use. The actual BIOS setting may be named “I/OAT DMA Engine”.
- X2APIC must be enabled.

The following BIOS settings are recommended on a compute server hosting VSR VMs that demands high packet-per-second forwarding performance.

- Hardware prefetching should be disabled.
- Adjacent cache line prefetching should be disabled.
- PCIe Active State Power Management (ASPM) should be disabled. (On Linux KVM hosts, ASPM can also be disabled by a kernel boot parameter.)
- ACPI C states should be disabled.
- Power management should be set to maximum or high performance.
- Turbo boost should be enabled.
- NUMA node interleaving should be disabled.

Note: Nokia recommends legacy boot mode instead of UEFI boot mode when the compute server is supporting VSR virtual machines.

11.1.7 Hypervisor and Host OS

VSR VMs can be deployed on compute hosts using either the KVM hypervisor or the VMware ESXi hypervisor that is part of the VMware vSphere suite. All virtual machines (VNF-Cs) comprising one VNF network element should use the same hypervisor technology.

11.1.8 KVM/QEMU Compute Hosts

A VSR VM can be deployed on a KVM compute host using a QCOW2 image available from Nokia. The VSR VM can be instantiated using libvirt tools (for example, virsh commands) or OpenStack. For more information on OpenStack, see the OpenStack Support section.

The KVM-QEMU hypervisor requires a Linux host OS. The host OS versions currently supported are:

- CentOS 7.0-1406 with 3.10.0-123 kernel
- CentOS 7.2-1511 with 3.10.0-327 kernel
- CentOS 7.4-1708 with 3.10.0-693 kernel
- Centos 7.5-1804 with 3.10.0-862 kernel
- Red Hat Enterprise Linux 7.1 with 3.10.0-229 kernel
- Red Hat Enterprise Linux 7.2 with 3.10.0-327 kernel
- Red Hat Enterprise Linux 7.4 with 3.10.0-693 kernel
- Red Hat Enterprise Linux 7.5 with 3.10.0-862 kernel
- Ubuntu 14.04 LTS with 3.13 kernel
- Ubuntu 16.04 LTS with 4.4 kernel

The QEMU package that comes standard with the Linux distribution should be used in most cases.

Note: Nokia recommends installing the QEMU Enterprise Virtualization packages of the CentOS Virtualization SIG (`yum install centos-release-qemu-ev qemu-kvm-ev`) on CentOS 7.X compute hosts that must support OpenStack-managed VMs with pinned vCPUs.

11.1.8.1 Linux Kernel Settings

On a KVM-based compute server hosting any type of VSR virtual machine, regardless of the VSR application or performance requirements, the following kernel boot settings are mandatory.

- Explicit (not transparent) huge pages that are 1GB in size must be used to back the guest memory of all VSR VMs and should be reserved at boot time. This is in addition to any huge page requirements for OVS-DPDK (if applicable).
- Intel IOMMU must be enabled to use SR-IOV or PCI pass-through.

- To avoid potential problems with spin lock loops, the Pause Loop Exiting gap must be set to 0.
- Security Enhanced Linux (SELinux) must be disabled or set to permissive mode.

11.1.8.2 Linux vSwitch Implementations

A virtual switch (vSwitch) is a software implementation of a Layer-2 bridge or Layer 2-3 switch in the host OS software stack. A vSwitch can allow VMs in the same host to exchange traffic without relying on switch hardware. A vSwitch can also allow multiple vNIC interfaces to share the same uplink into the data center.

The vSwitch implementations available to KVM-based compute hosts that are supported by VSR are:

- Linux bridge (vhost-net)
- Open vSwitch 2.3.0 (vhost-net)
- Open vSwitch 2.4.0 with DPDK 2.1.0 (vhost-user)
- Open vSwitch 2.5.0 with DPDK 2.2.0 (vhost-user). Requires QEMU 2.5.0 or later.

To connect a VSR virtual interface to a vSwitch in a KVM-based compute host, the hypervisor must provide a VirtIO back-end driver for the virtual interface. The back-end driver must be **vhost-net** or **vhost-user**. The E1000 driver is not supported.

Note: Nokia recommends using SR-IOV or PCI pass-through for virtual NIC interfaces intended to handle a high packet rate.

11.1.9 VMware ESXi Compute Hosts

A VSR-I VM supporting any application can be deployed on a VMware ESXi compute host using VMware vCloud Director (vCD) or the vSphere Web Client interface to a vSphere vCenter Server. Release 19.10.R6 supports the following solution sets:

- ESXi 6.0 Update 2, vCD 8.10 and vCenter Server 6.0 (vCloud NFV 1.5)
- ESXi 6.5 Update 1, vCD 8.20 and vCenter Server 6.5 (vCloud NFV 2.0)
- ESXi 6.7 and vCenter Server 6.7

Also, for the RR application only, a VSR-I virtual machine can also be deployed on a VMware ESXi 5.5 compute host. In this case, only the vSphere Web Client interface is supported.

11.1.9.1 vSphere Features

The following vSphere features are supported with VSR-I virtual machines, regardless of application or workload:

- Distributed Resource Scheduler (DRS) – but not fully-automated mode
- High Availability
- vSphere standard switch – connected to the guest using an E1000 or VMXNET3 driver
- vSphere distributed switch (vDS) – connected to the guest using an E1000 or VMXNET3 driver
- SR-IOV and PCI pass-through (NIC model dependent)

The following vSphere features are unsupported in Release 19.10.R6:

- DRS fully-automated mode
- vMotion
- Storage vMotion
- Fault Tolerance

11.1.9.2 NIC Settings

When using the VMXNET3 driver in the VSR software to connect its vNIC interfaces to a vSphere standard or distributed switch, then certain ESXi host-level settings are recommended to achieve optimal performance and avoid unnecessary packet drops. These best practice recommendations, which depend on the ESXi version, are provided in the *VSR Setup and Installation Guide*.

11.1.10 OpenStack Support

VSR virtual machines can be deployed on KVM compute hosts that are managed by OpenStack. The following OpenStack distributions and versions are supported:

- RDO OpenStack ‘Liberty’

- RDO OpenStack ‘Mitaka’
- RDO OpenStack ‘Newton’
- RDO OpenStack ‘Ocata’
- RDO OpenStack 'Pike'
- Red Hat OpenStack Platform 8 (OSP8)
- Red Hat OpenStack Platform 9 (OSP9)
- Red Hat OpenStack Platform 10 (OSP10)
- Red Hat OpenStack Platform 11 (OSP11)
- Red Hat OpenStack Platform 12 (OSP12)
- Red Hat OpenStack Platform 13 (OSP13)
- Mirantis OpenStack 9.0

When OpenStack is used to deploy VSR virtual machines on KVM compute, those compute hosts should be configured as follows:

- On each compute host the etc/nova/nova.conf configuration file must have the following lines to prevent oversubscription of resources:
  - cpu_allocation_ratio = 1.0
  - disk_allocation_ratio = 1.0
  - ram_allocation_ratio = 1.0
- On each compute host, the etc/nova/nova.conf configuration file must have a value for **vcpu_pin_set** that reserves a list of CPUs for the guests. The list of CPUs should be isolated from the host as described in the CPU Isolation section.
- On the OpenStack controller the etc/nova/nova.conf configuration file should be edited so that **scheduler_default_filters** lists both *AggregateInstanceExtraSpecFilter* and *NUMATopologyFilter*. The latter ensures that NUMA aware scheduling rules are applied.

11.1.11 VSR Lifecycle Management using CloudBand Application Manager (CBAM)

Lifecycle management of VSR-I instances running Release 16.0.R4, 19.5.R1, or higher is supported using CBAM 18.5. Lifecycle management of VSR-I instances running older software is supported using CBAM 17.5 and 17.5 SP1. VSR Release 19.5.R1 supports lifecycle management using CBIS 19. To on-board a VSR-I into CBAM, the official VNF package for VSR-I is required. The latest VNF package (with the filename vSRCbamSriovPackage_v19.5.1.1.zip) is downloadable from Nokia OLCS and contains all components needed for VSR-I instantiation, including:

- VNF descriptor (VNFD)
- VNFD metadata
- OpenStack HEAT templates (HOT)
- Mistral workflows
- Ansible playbooks
- Javascript helpers

Using the vSRCbamSriovPackage_19.5.1.zip package comes with the following restrictions:

- CBAM 18.5 is required.
- only grant-less operations are supported.
- OpenStack VMware is supported as a VIM.
- I/O technology supported for IOM-v ports are:
    - SR-IOV
    - PCI-PT
    - OVS
    - VirtIO
    - VRS.

NSP NFM-P, the element management system (EMS) of the VSR, supports integration with CBAM according to the ETSI MANO Ve-Vnfm-Em reference point. This allows NFM-P to actively monitor the status of virtual resources and to trigger certain LCM actions. Supported/tested combinations of VSR software, CBAM software, OpenStack VIM and NSP NFM-P software are listed in Table 28.

Table 28 **Supported Software Combinations for LCM of VSR-I**

SR OS software version	CBAM software version	OpenStack VIM	NFM-P software version
15.0.R4 and higher	17.5	Mitaka	17.9
15.0.R6 and higher	17.5 SP1	Mitaka	17.9
16.0.R1 and higher	17.5 SP1	Mitaka	18.6
16.0.R4 and higher	18.5	Liberty, Mitaka,Newton, Ocata, or Pike	18.12
16.0.R6 and higher/ 19.5.R1 and higher	18.5	Liberty, Mitaka, Newton, Ocata, or Pike	19.6

NFM-P is not mandatory for LCM of VSR-I instances. Table 29 summarizes the LCM actions that supported for VSR-I and describes how those actions are enhanced by NFM-P.

Table 29 **LCM Actions**

LCM Operation	Description
On-board VNF	From CBAM GUI or REST API, on-boards the VNF package in CBAM so that it is available for later use in VNF creation.
Create VNF	From CBAM GUI or REST API, this operation creates a uniquely identified VNF in CBAM.
Modify VNF information	From CBAM GUI or REST API, this operation modifies the VNF metadata and extensions from their defaults as defined in the VNFD.
Upgrade VNF package	From CBAM REST API, this operation upgrades the VNF package, not the VSR-I instance. Through this operation, a new VNFD containing corrections may be on-boarded.
Instantiate VNF	From CBAM GUI or REST API, this operation instantiates VSR-I on a specified OpenStack-based NFVI in grantless mode (with no NFVO integration). The necessary parameters are provided to CBAM in the form of a JSON file. The initial configuration of the newly instantiated VSR-I depends on the Ansible version used with CBAM. Further details are provided below. If NSP is integrated, CBAM notifies NFM-P upon successful creation and NFM-P auto-discovers the new instance. Note that when the VSR-I license is missing or invalid, the **configure** commands are skipped during instantiation.

Table 29 **LCM Actions (Continued)**

LCM Operation	Description
Terminate VNF	From CBAM GUI or REST API, this operation terminates a VSR-I instance.
Manual Heal VNFC	If NSP is integrated, manual heal request of VSR-I can be initiated by NFM-P; CBAM then implements the workflow by rebooting (first step) or rebuilding (second step) the VM and NSP is notified of the status. If NSP is not integrated, then the manual heal workflow must be initiated through the CBAM GUI or REST interface. Either VNFC name, combined slot (A,1), or CBAM machine Id (VSR_I) can be provided as input. During healing, soft/hard reboot or VM rebuild can be executed.
Auto-heal VSR-I	Requires NSP integration. If VSR-I becomes unreachable to NFM-P (according to user-defined policy rule), NSP raises an alarm and sends a heal workflow request to CBAM.

11.1.11.1 Initial Commissioning with CBAM 17.5

CBAM 17.5 is pre-packaged with Ansible 2.1. This version of Ansible does not support SR OS CLI, so the Ansible playbook provided in the VSR-I template package cannot perform any initial commissioning tasks. When using this version of CBAM, Nokia recommends modifying the **config.cfg** and **bof.cfg** files used to create the VSR-I instance in advance of deployment.

11.1.11.2 Initial Commissioning with CBAM 17.5 SP1

CBAM 17.5 SP1 includes both Ansible 2.1 and Ansible 2.3. Since Ansible 2.3 supports SR OS, this version should be used when deploying VSR virtual machines. This allows the Ansible playbook provided in the VSR-I template package to perform the following initial commissioning tasks (after it has determined that the VSR-I is alive by trying to connect to TCP port 22/SSH):

- Set the maximum snmp packet-size to 9216.
- Enable Telnet access.
- Enable FTP access.
- Enable SNMP.
- Set SNMP community string to allow SNMP read-write access by NSP.
- Configure the system IP address.
- Enable BOF persistence option.

11.2 VSR Deployment Considerations

11.2.1 VSR-a Appliances

VSR as an appliance (VSR-a) uses the integrated mode of VSR pre-integrated with RHEL and KVM on a Nokia AirFrame server. The VSR-a uses a single VSR-I VM with a pre-determined configuration of the system. In Release 19.10.R6, there are four models available for VSR-a: CN7, CN8, SN7, and SN8.

11.2.2 VSR-a Application Support

In Release 19.10.R6, the applications and licenses in Table 30 are approved for use with the VSR-a platform.

Table 30 **Applications and Licenses for VSR-a**

VSR Function	Release	VSR-a CN7	VSR-a CN8 [2]	VSR-a SN7	VSR-a SN8 [2]
Route Reflector (RR)	14.0.R6	✓	--	✓	--
Map-T Border Relay (Map-T)	15.0.R4	--	✓	--	✓
Network Group Encryption (NGE)	15.0.R6	--	✓	--	✓
Provider Edge (PE)	15.0.R6	✓	✓	✓	✓
Application Assurance (AA)	15.0.R6	--	✓ [1]	--	✓
Network Address Translation (NAT)	16.0.R1	--	✓ [1]	--	✓
Residential Gateway (RGW)	16.0.R7	--	--	--	✓
Broadband Network Gateway (BNG)	16.0.R7	--	--	--	✓
Wireless Lan Gateway (WLAN-GW)	16.0.R7	--	--	--	✓
Security Gateway (IPsec)	19.10.R1	--	--	--	✓

Notes:

1. CN8 will require the addition of an ISA card to the instantiation script to allow the configuration of additional applications.
2. Application combinations are also acceptable similarly to VSR-allowable combinations, such as WLAN-GW with NAT and AA. Performance and scale of various combinations may vary depending on the services configured.

11.2.3 VSR-a Supported Optics

This section describes the optics that are approved for use with the VSR-a platform.

Table 31 lists the GigE Optics Modules (SFP).

Table 31 **GigE Optics Modules (SFP)**

Nokia Part #	Short Description	Long Description
3HE00062CB	SFP - GIGE BASE-T RJ45 R6/6 DDM -40/85C	1-port 10/100/1000BASE-TX Small Form-Factor Pluggable (SFP) Copper Module, Cat5, RJ45 Connector, RoHS 6/6 compliant, Extended Temperature -40/85C
3HE05163AA	SFP - 100/1000 EX SGMII 40KM ROHS 6/6	1-port 100/1000Base Dual Rate SGMII Small Form-Factor Pluggable (SFP) Optics Module, 1310nm, 40km, DDM Compliant
3HE05164AA	SFP - 100/1000 LX SGMII 10KM DDM R6/6	1-port 100/1000Base Dual Rate SGMII Small Form-Factor Pluggable (SFP) Optics Module, 1310nm, 10km, DDM Compliant
3HE11904AA	SFP - GIGE BASE-T RJ45 R6/6 DDM -40/85C	1-port 10/100/1000BASE-TX Small Form-Factor Pluggable (SFP) Copper Module, Cat5, RJ45 Connector, RoHS 6/6 compliant, Extended Temperature -40/85C

Table 32 lists the 10GigE Optics Modules (SFP+).

Table 32 **10GigE Optics Modules (SFP+)**

Nokia Part #	Description	
3HE04823AA	SFP+ 10GE LR - LC ROHS6/6 0/70C	1-port 10GBASE-LR Small Form-Factor Pluggable+ (SFP+) Optics Module, Single Mode Fiber (SMF), 10km, 1310 nm, LC Connector, Digital Diagnostic Monitor (DDM), RoHS 6/6 compliant

Table 32 **10GigE Optics Modules (SFP+) (Continued)**

Nokia Part #	Description	
3HE04824AA	SFP+ 10GE SR - LC ROHS6/6 0/70C	1-port 10GBASE-SR Small Form-Factor Pluggable+ (SFP+) Optics Module, 850 nm, 26 to 300 meters, LC Connector, RoHS 6/6 compliant

Table 33 lists the 100GE Optics Modules (QSFP28).

Table 33 **100GE Optics Modules (QSFP28)**

Nokia Part #	Description	
3HE10550AA	QSFP28- 100G LR4 10KM LC ROHS6/6 0/70C	1-port 100GBase LR4 QSFP28 Optics Module, 10km, LC Connector, RoHS 6/6 compliant, Digital Diagnostic Monitor (DDM), 0/70C
3HE10551AA	QSFP28- 100G SR4 100m, MPO ROHS6/6 0/70C	1-port 100GBase SR4 QSFP28 Optics Module, 100m, MMF, MPO Connector, RoHS 6/6 compliant, Digital Diagnostic Monitor (DDM), 0/70C
3HE10552AA	QSFP28- 100G CWDM4 2KM LC ROHS6/6 0/70C	1-port 100GBase CWDM4 QSFP28 Optics Module, 2km, LC Connector, RoHS 6/6 compliant, Digital Diagnostic Monitor (DDM), 0/70C

Only a single QSFP28-LR4 optic can be used in the 100GigE NIC at a time.

If SFP/SFP+ optics are to be used in a QSFP28 port, then the use of a QSFP-to-SFP+ adapter is required. Nokia recommends the Mellanox MAM1A00A-QSA adapter.

11.2.4 VM Memory Requirements by Function Mix

Table 34 shows the minimum memory for VSR-I by function mix.

Table 34 **Minimum Memory for VSR-I by Function Mix**

Functions	Minimum Memory for VSR-I VM (GB)
PE	4
RR	4
NAT	28 (for profile-1) 56 (for profile-2)
Transit AA (Res 8K mode)	8
PE + AA (VPN 1K mode)	8
IPsec	20 [1]
IPSec + AA firewall	30 [2]
LNS	24
BNG without ISA	16 32 (>32k queues)
BNG with BB-ISA	24 48 (>32k queues)
vRGW/WLAN-GW	32
vRGW/WLAN-GW with AA	40

Notes:

1. Requirement to achieve maximum number of IPsec tunnel scale.
2. Requirement to support AA at maximum number of IPsec tunnel scale.

11.3 Supported VSR Features

This section provides a summary view of SR OS features supported on VSR including features introduced in previous releases that are supported on VSR.

11.3.1 New Features Supported in Release 19.10.R6

- SR-IOV support on the Intel XXV710-DA2 NIC

11.3.2 New Features Supported in Release 19.10.R5

No new features supported on VSR are introduced in Release 19.10.R5.

11.3.3 New Features Supported in Release 19.10.R4

No new features supported on VSR are introduced in Release 19.10.R4.

11.3.4 New Features Supported in Release 19.10.R3

No new features supported on VSR are introduced in Release 19.10.R3.

11.3.5 New Features Supported in Release 19.10.R2

No new features supported on VSR are introduced in Release 19.10.R2.

11.3.6 New Features Supported in Release 19.10.R1

- NGE for L2 auto-gre-sdp Services
- Unnumbered Subscriber Interfaces
- EHS/CRON: allow scripts to execute bypassing MD interface locks
- YANG model “must” statements (subsequently removed in Release 19.10.R3)
- MD-CLI **insert** command for user-ordered lists
- Complete support for source/target combinations in NETCONF <copy-config> operations
- Support multiple error messages with NETCONF RPCs
- Local command accounting log events for NETCONF
- Support port 22 in addition to 830 for NETCONF
- LI support in MD-CLI and NETCONF
- NETCONF insert command for user-ordered lists
- Bytes encoding for gNMI subscribe RPCs
- Local command accounting log events for gRPC

- gNOI Certificate Management
- MD-CLI configuration output in JSON format
- MD-CLI pwc Command Enhancements
- MD-CLI **ping** command
- MD-CLI **traceroute** command
- MD-CLI admin tree
- Accounting Statistics Alignment
- New Counter for Oversized Packets
- BGP Remove-Private Replace Option
- BGP Segment Routing Using the Prefix SID Attribute
- Static and BGP SR Policies with IPv6 Endpoint
- BGP ORR Enhancement
- BGP Unresolved Route Leaking from Base to VPRN
- RFC 6549 OSPFv2 Multi-Instance Extensions
- Support for Strict IS-IS Weighted ECMP
- Route Policy Support of Named Entries
- Route Policy Action to Suppress BGP Route Installation
- Multi-Chassis Synchronization of RADIUS Usage Counters
- Diameter Multi-Chassis Redundancy on New Base
- Data-Triggered ESM Host Mobility
- Sub-profile and sla-profile string length increase (16B->32B)
- L2-aware access mode (Bridged home with nh-mac antispoof)
- Flow Attributes
- Parental Control with Rest-API
- GRE Fragmentation and Reassembly
- GRE Termination on Interface IP Address

11.3.7 Supported Features

11.3.7.1 L1/L2 Networking

- Ethernet ports
- Link aggregation groups (LAG)

- LACP
- MC-LAG
- 802.1Q VLANs
- QinQ
- Jumbo frames [1]
- Interface statistics
- Network port
- Hybrid port
- Access port
- Spoke-SDP IP interfaces
- Port cross-connect/PXC [2]

Notes:

1. The maximum MTU size (at the VSR vNIC port) depends on the deployment model. For VSR-I, the maximum MTU is 9212 bytes. Due to host dependencies, it may not be possible to reach these limits on vNIC ports using VirtIO.
2. PXC FPE extensions are not supported.

11.3.7.2 IPv4 and IPv6 Routing Protocols and Scaling

- Static routes
- OSPFv2, OSPFv3, IS-IS, RIP, and RIPng
- BGP4 and MP-BGP
- ARP, IPv6 ND
- ICMP, ICMPv6
- Equal-cost multipath (ECMP)
- Unequal-cost multipath/weighted ECMP (for BGP IP routes and IGP shortcuts over RSVP-TE tunnels)
- Unicast RPF
- VRRP
- BGP FlowSpec (functional parity)
- Default EBGP route propagation behavior without policies
- RFC 5549 support for IPv4 BGP Routes
- Multiple peer AS Support for dynamic BGP sessions
- IS-IS reference bandwidth increase (to support 64 x 100 GE LAG)

- OSPF reference bandwidth increase
- Increased interface description from 160 to 256 characters

11.3.7.3 IPv4 and IPv6 Multicast Protocols and Scaling

- Base router and VPRN support for the following protocols:
  - IGMP v1/v2/v3
  - MLDv1/v2
  - PIM
  - MSDP
- Up to 64K IPv4/IPv6 multicast groups per system
- BGP multipath enhancements
- IPv4/IPv6 multicast protocols in Base router and VPRNs (IGMP, MLD, PIM, MSDP)

11.3.7.4 ECMP/LAG Hashing

- Traffic load-balancing based on:
  - L3 and L4 fields
  - L3-only fields
  - source-only fields
  - destination-only fields
- ECMP/LAG hash:
  - configurable include/exclude L4 fields
  - load-balancing based on source fields only or destination fields only
- Hash label:
  - VPLS and epipe service contexts

11.3.7.5 OAM

- BFD (centralized and distributed)
- SDP ping
- VXLAN ping

11.3.7.6 IGP

- OSPF External Type 1 Overload
- IS-IS Minimum Remaining Lifetime
- RFC 6549 OSPFv2 Multi-Instance Extensions
- Support for Strict IS-IS Weighted ECMP

11.3.7.7 Unicast Routing

- BGP Default Route Origination
- BGP AS-Override in Base Context
- BGP Convergence - Delayed Route Advertisement
- Improved Handling of BGP MED in Route Policy Actions
- RPKI on VPRN BGP PE-CE session

11.3.7.8 MPLS and Segment Routing

- LDP for IPv4 FECs
- RSVP point-to-point LSPs
- MPLS RSVP Fast Reroute (FRR)
- LDP-over-RSVP
- BGP label-unicast IPv4 (BGP RFC 3107)
- 6PE
- OSPFv2/IS-IS shortcuts to IPv4 prefixes (using LDP or RSVP)
- BGP shortcuts to IPv4 prefixes (using LDP, RSVP, or BGP 3107)
- OSPFv2 segment routing extensions
- IS-IS segment routing extensions
- SR traffic engineering (SR-TE)
- BGP segment routing policies
- Segment-Routing TE using MPLS dataplane

11.3.7.9 Filtering, OpenFlow, Control Plane Protection

- Ingress IPv4 and IPv6 filters
- Egress IPv4 and IPv6 filters
- IP filter override for R-VPLS services
- All standard match criteria supported by SR OS
- Standard actions: Forward, drop and HTTP redirect
- Conditional actions: **drop-extracted-traffic** (for control plane protection), drop based on packet length, drop based on TTL
- Ingress PBR actions: forward to next-hop, forward to router (another routing instance), redirect-policy
- NAT action
- Reassembly action
- Filter logging (ingress and egress)
- Distributed CPU protection (static policers only)
- IPv4 BGP FlowSpec
- IPv6 BGP FlowSpec
- OpenFlow
- Filter Rate-Limit Action

11.3.7.10 Layer-2 VPNs and DCGW

- Ethernet VLL signaled by T-LDP using MPLS or GRE transport
- Ethernet VLL signaled by BGP using MPLS or provisioned GRE SDP transport
- Ethernet VLL using L2TPv3 (static)
- Ethernet VLL signaled by BGP-EVPN using MPLS
- Static Ethernet VLL using VXLAN IPv4 transport
- Ethernet VPLS signaled by T-LDP using MPLS or GRE transport
- Ethernet VPLS with BGP-AD signaled spoke-SDPs
- Ethernet VPLS signaled by BGP using MPLS or provisioned GRE SDP transport
- Ethernet VPLS signaled by BGP-EVPN using MPLS or VXLAN transport
- Virtualized DCGW with Nuage VSD integration, including the support for fully dynamic XMPP model
- Routed-VPLS
- Resiliency

  - Pseudowire redundancy
  - Dual-homed VPWS/VLL
  - BGP multi-homing for VPLS
  - MC-LAG
  - STP, RSTP, MSTP
- IGMP snooping in VPLS services (excluding BGP-VPLS, BGP-AD, EVPN and Routed-VPLS services)
- VPLS endpoint scale increase from 50 to 128

11.3.7.11 Layer-3 Services

- Internet access (IES services)
- RFC 4364 IPv4 VPNs using MPLS or GRE transport
- RFC 4659 IPv6 VPNs using MPLS or GRE transport
- IP VPN inter-AS option B
- IP-in-IP and GRE IP tunneling (using **isa-tunnel-v**)
- GRT lookup and VPRN-to-GRT route leaking
- VPRN and IES spoke-SDP IP interfaces
- SAP scale increase to 32k
- IP filtering on R-VPLS interfaces

11.3.7.12 Network Group Encryption (NGE)

- SDP encryption (MPLS and GRE)
- VPRN encryption (MPLS, GRE, MPLSoUDP)
- Router interface encryption
- NGE for L2 auto-gre-sdp Services

11.3.7.13 Quality of Service (QoS)

- Up to 14 GB of packet buffer memory per-VSR
- Ingress pre-classification for class-aware early discard (optional, enabled using the **config**>**system**>**congestion-management** command)

- Ingress classification to forwarding-class based on 802.1p, DSCP, MPLS EXP or IPv4/IPv6 filter rules
- Egress re-classification
- Ingress and egress policing/hierarchical policing
- Up to 64k egress policers per-VSR
- Egress marking of 802.1p, DSCP, or MPLS EXP
- Up to 128k egress queues per VSR
- Egress queue shaping based on configurable PIR and MBS
- Egress H-QoS with queue parenting to port or user scheduler
- Up to three (3) tiers of egress user schedulers
- Eight (8) strict priority levels per egress user scheduler
- Weighted round robin (WRR) scheduling in each scheduler level
- Aggregate SAP limit (for example, **config**>**service**>**epipe**>**sap**>**egress**>**agg-rate**), including frame-based accounting
- Aggregate subscriber rate limit (for example, **config**>**service**>**vprn**>**sub-if**>**grp-if**>**sap**>**egress**>**agg-rate**), including frame-based accounting
- **packet-byte-offset** configurable per-queue (Note that when both **packet-byte-offset** and **frame-based-accounting** are configured, they work together to determine the effective offset value.)

11.3.7.14 Mirroring and Lawful Intercept

- Ether and IP-only mirror types
- Debug mirror sources: ports
- LI mirror sources: subscribers, SAPs, spoke-SDPs
- Mirror destinations: SAPs, spoke-SDPs
- Basic LI management infrastructure
- Routable LI encap (IP/UDP and IP/GRE)
- Pre-NAT (private IP) and post-NAT (public IP) subscriber mirroring/LI (Note that pre-NAT mirroring/LI is only supported with L2-aware NAT.)
- Mirroring/LI into a spoke-SDP (originate and terminate)

11.3.7.15 Broadband Network Gateway (BNG)

- Routed-CO model of Enhanced Subscriber Management (ESM) on numbered subscriber interface and group interface
- Dual-Stack IPoE subscriber management
- Dual-Stack PPPoE sessions
- Static SAP and MSAP
- 1:1 and N:1 VLANs
- Managed routes (IPv4/IPv6): RADIUS and BGP (no RIP, no PPPoE IPCP subnet negotiation)
- Subscriber authentication using LUDB, RADIUS, Diameter NASREQ/Gx (no support for flow-based PCC rules)
- Dynamic QoS overrides
- Dynamic filter overrides
- RADIUS Accounting per session/host/SPI
- LAG for subscriber access
- HTTP Redirect
- H-QoS for subscriber hosts (support for single user scheduler per subscriber)
- Data-triggered SAPs and ESM hosts (support for stateless redundancy, no support for stateful redundancy)
- L2TP LAC/LTS/LNS (no MLPPP)
- Credit control using category-maps: RADIUS Credit Control, Diameter Gy, Diameter Gx Usage Monitoring, Idle-Timeout.
- Wholesale/retail (no **ipoe private-retail-subnets**)
- Subscriber interface 128-bit WAN mode
- L2TP-tunnel accounting
- DHCPv6-PD as managed route
- Python for following applications: DHCPv4, DHCPv6, DTC/ESM, RADIUS (no support for extended VSAs), Python cache
- SRRP with MCS on VSR-BNG for following applications: **local-dhcp-server**, **sub-mgmt-ipoe**, **sub-mgmt-pppoe**, and **python**
- VSR-BNG scale increase to 64k subscribers with 2 egress queues and 1 ingress policer per subscriber
- ESM-over-GTP support for fixed wireless access
- Increase total length of Alc-ToServer-Dhcp-Options and Alc-ToServer-Dhcp6-Options attributes to 1000 bytes
- Unnumbered subscriber interfaces

- Multi-Chassis Synchronization of RADIUS Usage Counters
- Diameter Multi-Chassis Redundancy on New Base
- Data-Triggered ESM Host Mobility
- Sub-profile and sla-profile string length increase (16B->32B)

Note: LNS features require the **isa-bb-v** to be configured.

11.3.7.16 Network Address Translation (NAT)

- LSN44 (deterministic and non-deterministic)
- NAT64
- DS-Lite
- L2-Aware NAT in vBNG context
- UPnP for L2-Aware NAT
- Geo-redundancy
- Per-host DNS Override Using Destination NAT
- NAT scale increase using profiles
- L2-aware access mode (Bridged home with nh-mac antispoof)

Note: NAT features require the **isa-bb-v** to be configured.

11.3.7.17 MAP-T Border Relay (MAP-T BR)

- Hub-and-Spoke model
- Full routing support (IS-IS, BGP, OSPF, RIP)
- Upstream MAP-T anti-spoof
- Multiple MAP-T domains in the same routing context
- Upstream/downstream fragmentation
- MSS adjust and MTU support per domain
- Forward/drop statistics collection per domain
- Fragmentation statistics

- Configuration logging

11.3.7.18 IPsec Security Gateway (SeGW)

- IKEv1 static/dynamic LAN-to-LAN tunnel with pre-shared key authentication
- IKEv1 remote-access tunnel with plain-xauth-psk authentication
- IKEv2 static/dynamic LAN-to-LAN tunnel and remote-access tunnel
- IKEv2 tunnel authentication method: **psk**/**psk-radius**/**cert-auth**/**cert-radius**/**eap**/**auto-eap**/**auto-eap-radius**
- IKEv2 remote-access tunnel internal address assignment methods: RADIUS/local address pool/external DHCPv4/v6 server
- Encryption Algorithm: DES/3DES/AES-CBC/AES-GCM/AES-GMAC
- Authentication Algorithm: MD5/SHA1/SHA256/SHA384/SHA512
- Diffie-Hellman Group: 1/2/5/14/15/19/20/21
- Perfect Forward Secrecy
- NAT-T support
- IPv4 and IPv6 support
- Multi-chassis-IPsec
- IKEv2 Fragmentation
- Client Lockout
- Auto CRL Update
- TCP MSS Adjust
- IPsec private side fragmentation
- IPsec certificate authentication support certificate chain with 3 or more CAs
- OCSP
- New crypto algorithms for IPsec:
  - ECDSA
  - RFC 7427
- IPsec Secure Interface
- IPsec-fastpath External Offload
- Ignore ANY TS for RRI creation

Note: IPsec features require the **isa-tunnel-v** to be configured.

11.3.7.19 Application Assurance (AA)

- AA use cases:
  - **dns-ip-cache**
  - AA policers (all types)
  - AA **http-redirect**
  - **http-notification** (IBN)
  - **http-enrich**
  - **url-filter** using the ICAP interface
  - **url-filter url-list** for local URL filtering
- AA Group Partition features:
  - AA RADIUS accounting
  - AA **event-log** (**syslog** export)
  - Stateful Firewall (FW) (**gtp**, **gtp-filter**, **sctp-filter**, **session-filter**)
  - AA policy application/AG/ASO/**app-filter**, signatures
  - AA policy charging-groups
  - AA **transit-ip-policy** (IPv4/IPv6)
  - AA transit prefix lists
- Statistics - all AA XML and IPFIX statistics records export
- AA Hi/Lo resource watermark alarms
- AARP - Local protection
- AA support on the following services:
  - Epipe SAP and spoke-SDP (MPLS+GRE)
  - VPLS SAP and spoke-SDP
  - IES/VPRN SAP and spoke-SDP
  - IES/VPRN ESM (in VSR-BNG, VSR-WLANGW, or VSR-vRGW with AA)
  - IES/VPRN DSM (in VSR-WLANGW with AA)
  - IES/VPRN IPsec Private SAP (in VSR-PE or VSR-SeGW)
  - AA tunnel support (DS-Lite, 6RD/6to4, Teredo, GRE)
- AA signatures configuration-only upgrade without VSR system upgrade
- AA vRGW Nested-Router Detection
- Configurable Cflowd parameters
- Flow Attributes
- Parental Control with Rest-API

Note: AA features require the **isa-aa-v** to be configured.

11.3.7.20 Virtual Residential Gateway (vRGW) and Wireless-LAN Gateway (WLAN-GW)

vRGW and WLAN-GW are supported, with the exception of the items listed in SR OS Features not Supported on VSR.

- NGE Support for WLAN-GW L2oMPLSoGRE
- Home LAN Extension: Policers

11.3.7.21 VSR Platform Infrastructure

- Allow VSR on host with NIC trust mode off - v4 traffic only
- Red Hat OpenStack Platform 11.0 (Ocata) support
- VSR on Centos 7.4 host running KVM
- VSR on Red Hat Enterprise Linux 7.4 host running KVM
- VSR on VMware ESXi 6.5 host
- Intel Skylake CPU support (KVM and VMware)
- Intel X710 4x10G NIC support (KVM only)
- SR-IOV support for Mellanox Connect X4 NICs on VMware hosts
- Mellanox Connect X5 NIC (MCX-516A) support (KVM and VMware)
- Optimized hyper-threading support on KVM and VMware hosts
- Support for up to 56 vCPUs (on one NUMA node) per VSR virtual machine
- VSR-I software license expiry change without reboot
- OpenStack Mirantis 9.0 support
- Intel X710 NIC support for untagged multicast frames in SR-IOV mode (Centos/Red Hat 7.4 or higher)
- VSR on VMware ESXi 6.7 host
- Red Hat OpenStack Platform 12.0 (Pike) support
- VSR on Centos 7.5/RHEL 7.5 host running KVM
- VSR lifecycle management using CBAM 18.5
- Intel X710 NIC support on KVM using 2.14.1-k driver

- Intel X722 NIC support on VSR-a (PCI pass-through)
- HPE 2-port 10G and 25G NICs (SR-IOV and PCI-passthrough)
- It is no longer required to blacklist the SR-IOV i40evf driver
- GRE Fragmentation and Reassembly
- GRE Termination on Interface IP Address
- SR-IOV support on the Intel XXV710-DA2 NIC

11.3.7.22 Flex PW-Ports

- MPLS SDP binding
- L2oGRE using IPv4 or IPv6 transport
- EVPN VXLAN v4
- Static VXLAN v4
- EVPN-VPWS MPLS
- EVPN-VPWS MPLSoUDP

11.3.7.23 Model-Driven Management

- Configuration using model-driven (MD) interfaces (the MD-CLI, NETCONF, gRPC) – Nokia SR OS YANG model coverage is equivalent to hardware routers, for functionality that is supported by VSR. Contact a Nokia representative for information about the availability of specific configuration paths for model-driven management.
- State information retrieval using model-driven (MD) interfaces (the MD-CLI, NETCONF, gRPC) – Nokia SR OS YANG model coverage is equivalent to hardware routers, for functionality that is supported by VSR. Contact a Nokia representative for information about the availability of specific state paths for model-driven management.
- Telemetry using the gNMI Subscribe RPC, supporting the following modes:
  - ONCE
  - SAMPLE
  - ON_CHANGE
  - TARGET_DEFINED
- gNMI Telemetry
- Model driven management infrastructure support (the MD-CLI, NETCONF, gRPC) on VSR equivalent to hardware 7x50 routers

- MD-CLI **insert** command for user-ordered lists
- EHS/CRON: allow scripts to execute bypassing MD interface locks
- YANG model “must” statements (subsequently removed in Release 19.10.R3)

11.3.7.24 NETCONF

- Complete the support for source/target combinations
- Support multiple error messages with NETCONF RPCs
- Local command accounting log events for NETCONF
- Support port 22 in addition to 830 for NETCONF
- LI Support in the MD-CLI and NETCONF
- NETCONF insert command for user-ordered list

11.3.7.25 Telemetry

- Bytes encoding in telemetry notifications

11.3.7.26 gNMI and Other gRPC Items

- Local command accounting log events for gRPC
- gNOI Certificate Management

11.3.7.27 MD-CLI

- MD-CLI configuration output in JSON format
- MD-CLI pwc Command Enhancements
- MD-CLI **ping** command
- MD-CLI **traceroute** command
- MD-CLI admin tree

11.3.7.28 System

- Accounting Statistics Alignment

11.3.7.29 Platform: PORT

- New Counter for Oversized Packets

11.3.7.30 BGP

- BGP Remove-Private Replace Option
- BGP Segment Routing Using the Prefix SID Attribute
- Static and BGP SR Policies with IPv6 Endpoint
- BGP ORR Enhancement
- BGP Unresolved Route Leaking from Base to VPRN

11.3.7.31 Routing Infrastructure

- Route Policy Support of Named Entries
- Route Policy Action to Suppress BGP Route Installation

11.4 SR OS Features not Supported on VSR

Table 35 shows SR OS features that are supported on SR OS hardware routers but not on VSR. If you have concerns or doubts about VSR support for a particular feature, please contact your Nokia representative. Unsupported features in the following table reference classic CLI paths. The corresponding model-driven paths are also unsupported.

Table 35 **SR OS Features not Supported on VSR**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
AA datapath CPU watermarks	config>app-assure>datapath-cpu-high-wmark config>app-assure>datapath-cpu-low-wmark	N
AA debug mirror	debug>mirror-source> isa-aa-group	N
AA IES/VPRN SAP transit prefix	config>isa>aa-group>transit-prefix-ipv4-entries config>isa>aa-group>transit-prefix-ipv6-entries config>app-assure>group>transit-prefix-policy	N
AA inter-chassis AARP	config>app-assure>aarp>peer, peer-endpoint, master-selection-mode config>service>ies>aarp-interface config>service>vprn>aarp-interface	N
AA overload cut through	config>isa>aa-group>isa-overload-cut-through	N
AA signatures SW upgrade without VSR system upgrade	--	N
APS	config>port>aps	Y
ATM	config>qos>atm-td-profile config>service>apipe config>service>epipe>sap>atm config>service>ies>if>sap>atm config>service>ies>sub-if>grp-if>sap>atm config>service>vpls>sap>atm config>service>vprn>if>sap>atm config>service>vprn> sub-if>grp-if>sap>atm config>subscr-mgmt>msap-policy>atm config>system>atm	Y
BFD on CPM P-chip	config>router>if>bfd>type cpm-np config>router>if>ipv6>bfd>type cpm-np config>service>ies>if>bfd>type cpm-np config>service>ies>if>ipv6>bfd>type cpm-np config>service>vprn>if>bfd type cpm-np config>service>vprn>if>ipv6>bfd type cpm-np config>service>vprn>nw-if>bfd type cpm-np	Y
BGP EVPN multi-homing (ES end)	config>service>system>bgp-evpn>ad-per-es-route-target config>service>system>bgp-evpn>ethernet-segment config>service>system>bgp-evpn>route-distinguisher	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
BGP EVPN multi-homing (aliasing)	--	N
BGP EVPN - PBB	--	N
BGP EVPN over VXLAN - non-system-IP termination	--	N
BGP EVPN over VXLAN IPv6	--	N
BGP EVPN - Etree	--	N
BGP EVPN over MPLS - PIM snooping	--	N
BGP EVPN blackhole MAC	--	N
BGP FRR	config>router>bgp>backup-path config>service>vprn>bgp>backup-path config>service>vprn>enable-bgp-vpn-backup	Y
Call-trace	config>call-trace	N
Central frequency clock	config>system>sync-if-timing config>card>mda>sync-e config>port>ethernet>ssm	Partial
Card reset on recoverable error	config>card>reset-on-recoverable-error	Y
Cflowd	config>cflowd config>filter>ip-filter>entry>filter-sample config>filter>ip-filter>entry>interface-disable-sample config>filter>ipv6-filter>entry>filter-sample config>filter>ipv6-filter>entry>interface-disablesample config>router>if>cflowd-parameters	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
CIR/CBS for queues and schedulers	config>qos>port-scheduler-policy>group>percent-rate cir config>qos>port-scheduler-policy>group>rate cir config>qos>port-scheduler-policy>level cir config>qos>port-scheduler-policy>level percent-cir config>qos>port-scheduler-policy>orphan-override cir-weight cir-level config>qos>sap-egress>queue>adaptation-rule cir config>qos>sap-egress>queue>cbs config>qos>sap>egress>queue>parent cir-weight cir-level config>qos>sap>egress>queue>port-parent cir-weight cir-level config>qos>sap>egress>queue>rate cir config>qos>scheduler-policy>tier>scheduler>parent cir-weight cir-level config>qos>scheduler-policy>tier>scheduler>port-parent cir-weight cir-level	N
Circuit emulation	config>mirror>mirror-dest>sap>cem config>service>cpipe config>service>epipe>sap>cem	Y
Class-based forwarding	config>router>ldp>class-forwarding config>service>sdp>class-forwarding config>router>mpls>lsp>class-forwarding config>router>mpls>lsp-template>class-forwarding	N
CLI rollback	config>system>rollback	N
Connection profiles	config>connection-profile	Y
CPM filters	config>system>security>cpm-filter	Y
CPM queues	config>system>security>cpm-queue	Y

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
CPU protection	config>router>if>cpu-protection config>service>epipe>sap>cpu-protection config>service>epipe>spoke-sdp>cpu-protection config>service>ies>if>cpu-protection config>service>ies>if>sap>cpu-protection config>service>ies>if>spoke-sdp>cpu-protection config>service>ies>sub-if>grp-if>sap>cpu-protection config>service>template>vpls-sap-template>cpu-protection config>service>vpls>mesh-sdp>cpu-protection config>service>vpls>sap>cpu-protection config>service>vpls>spoke-sdp>cpu-protection config>service>vprn>if>cpu-protection config>service>vprn>if>sap>cpu-protection config>service>vprn>nw-if>cpu-protection config>service>vprn> sub-if>grp-if>sap>cpu-protection config>subscr-mgmt>msap-policy>cpu-protection config>system>security>cpu-protection	Y
Distributed CPU protection - dynamic and local-monitoring policers	config>card>fp>dist-cpu-protection config>sys>security>dist-cpu-protection>policy>local-monitoring-policer config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters	Partial
Distributed CPU protection - static policer exceed action low-priority	config>sys>security>dist-cpu-protection>policy>static>exceed-action low-priority	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Entropy label	config>router>entropy-label config>router>ldp>entropy-label-capability config>router>mpls>entropy-label config>router>mpls>lsp>entropy-label config>router>mpls>lsp-template>entropy-label config>router>rsvp>entropy-label-capability config>service>epipe>bgp-evpn>mpls>entropy-label config>service>epipe>spoke-sdp>entropy-label config>service>pw-template>entropy-label config>service>vpls>bgp-evpn>mpls>entropy-label config>service>vpls>mesh-sdp>entropy-label config>service>vpls>spoke-sdp>entropy-label config>service>vprn>entropy-label	N
ESM in VPLS service	--	N
ESM-over-GTP support for fixed wireless access	--	N
Ethernet CFM	config>eth-cfm config>lag>eth-cfm config>port>ethernet>eth-cfm config>service>epipe>eth-cfm config>service>epipe>sap>eth-cfm config>service>epipe>spoke-sdp>eth-cfm config>service>ies>eth-cfm config>service>ies>if>sap>eth-cfm config>service>ies>if>spoke-sdp>eth-cfm config>service>ies>sub-if>grp-if>sap>eth-cfm config>service>vpls>eth-cfm config>service>vpls>mesh-sdp>eth-cfm config>service>vpls>sap>eth-cfm config>service>vpls>spoke-sdp>eth-cfm config>service>vprn>eth-cfm config>service>vprn>if>sap>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm	Y
Ethernet rings	config>eth-ring	N
Ethernet satellites	config>system>satellite>eth-sat	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	**Associated CLI (Not Exhaustive)**	**Commands Removed**
Ethernet tunnels	config>eth-tunnel config>router>l2tp>eth-tunnel config>router>l2tp>group>eth-tunnel config>service>epipe>sap>eth-tunnel config>service>vpls>sap>eth-tunnel config>service>vprn>l2tp>eth-tunnel config>service>vprn>l2tp>group>eth-tunnel	Y
EXP secondary shaper	config>port>ethernet>egress>exp-secondary-shaper	Y
Frame relay	config>service>fpipe config>qos>mc-fr-profile-egress config>qos>mc-fr-profile-ingress	Partial
GMPLS and LMP	config>gmpls-tun-grp config>router>gmpls config>router>lmp	Y
Hash label	config>service>fpipe>spoke-sdp>hash-label config>service>ies>interface>spoke-sdp>hash-label config>service>ipipe>spoke-sdp>hash-label config>service>vprn>hash-label config>service>vprn>interface>spoke-sdp>hash-label	N

Table 35 SR OS Features not Supported on VSR (Continued)

Feature	Associated CLI (Not Exhaustive)	Commands Removed
HS-MDAv2	config>card>mda>egress>hsmda-agg-queue-burst config>port>ethernet>hsmda-scheduler-overrides config>qos>hsmda-pool-policy config>qos>hsmda-scheduler-policy config>qos>hsmda-slope-policy config>qos>hsmda-wrr-policy config>qos>network-queue>egress-hsmda config>qos>network-queue>fc>egress-hsmda config>qos>queue-group templates>egress>queue-group>hsmda-queues config>qos>sap-egress>fc>hsmda config>qos>sap-egress>hsmda-queues config>service>epipe>sap>egress>hsmda-queue-override config>service>ies>if>sap>egress>hsmda-queue-override config>service>vpls>sap>egress>hsmda-queue-override config>service>vprn>if>sap>egress>hsmda-queue-override config>subscr-mgmt>sub-profile>hsmda config>lag>port-type	Partial
Ipipe services	config>service>ipipe	N
Ingress queues and schedulers	config>qos>sap-ingress>queue config>qos>sap-ingress>fc>broadcast-queue config>qos>sap-ingress>fc>multicast-queue config>qos>sap-ingress>fc>queue config>qos>sap-ingress>fc>unknown-queue config>serice>epipe>sap>ingress>scheduler-policy config>service>ies>if>sap>ingress>scheduler-policy config>service>ies>sub-if>grpif>sap>ingress>scheduler-policy config>service>template>vpls-sap-template>ingress>scheduler-policy config>service>vpls>sap>ingress>scheduler-policy config>service>vprn>if>sap>ingress>scheduler-policy config>service>vprn>sub-if>grpif>sap>ingress>scheduler-policy config>subscr-mgmt>sla-profile>ingress>qos>queue config>subscr-mgmt>sub-profile>ingress>scheduler-policy	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
IP filter PBR actions	config>filter>ip-filter>entry>egress-pbr config>filter>ip-filter>entry>pbr-down-action-override config>filter>ip-filter>entry>sticky-dest config>filter>ip-filter>entry>action>extended-action config>filter>ip-filter>entry>action>forward esi config>filter>ip-filter>entry>action>forward lsp config>filter>ip-filter>entry>action>forward sap config>filter>ip-filter>entry>action>forward sdp config>filter>ip-filter>entry>action>forward vprn-target config>filter>ipv6-filter>entry>egress-pbr config>filter>ipv6-filter>entry>pbr-down-action-override config>filter>ipv6-filter>entry>sticky-dest config>filter>ipv6-filter>entry>action>extended-action config>filter>ipv6-filter>entry>action>forward esi config>filter>ipv6-filter>entry>action>forward lsp config>filter>ipv6-filter>entry>action>forward sap config>filter>ipv6-filter>entry>action>forward sdp config>filter>ipv6-filter>entry>action>forward vprn-target	N
IP filter remark action	config>filter>ip-filter>entry>action>remark config>filter>ip-filter>entry>action>extended-action remark config>filter>ipv6-filter>entry>action>remark config>filter>ipv6-filter>entry>action>extended-action remark	N
IP filter pattern match	config>filter>ip-filter>entry>action>drop pattern config>filter>ip-filter>entry>action>forward-when pattern config>filter>ipv6-filter>entry>action>drop pattern config>filter>ipv6-filter>entry>action>forward-when pattern config>filter>ip-filter>entry>action>rate-limit <value> pattern config>filter>ipv6-filter>entry>action>rate-limit <value> pattern	N
Filter type src-mac and packet-length	config>filter>ip-filter type src-mac config>filter>ip-filter type packet-length config>filter>ipv6-filter type src-mac config>filter>ipv6-filter type packet-length	N
IP interface ingress statistics on VPRN and IES SAPs	config>service>ies>if>enable-ingress-stats config>service>vprn>if>enable-ingress-stats	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
IP VPN Carrier Serving Carrier (CSC)	config>service>vprn>carrier-carrier-vpn config>service>vprn>nw-if	N
IP VPN inter-AS model C	--	N
IP VPN LPP/LPN	config>service>vprn>label-mode next-hop config>router>policy-options>policy-statement>entry>action>advertise-label per-prefix	N
IPv6: 6-to-4	--	N
ISSU	--	
L2TPv3	config>router>l2tp>group>l2tpv3 config>router>l2tp>l2tpv3 config>service>epipe>sap>l2tpv3-session config>service>ies>if>sap>l2tpv3-session config>service>vpls>sap>l2tpv3-session config>service>vprn>if>sap>l2tpv3-session config>service>vprn>l2tp>group>l2tpv3 config>service>vprn>l2tp>l2tpv3	Y
LAG link-map profiles	config>lag>link-map-profile config>router>if>lag-link-map-profile config>service>epipe>sap>lag-link-map-profile config>service>ies>if>sap>lag-link-map-profile config>service>ies>sub-if>grp-if>sap>lag-link-map-profile config>service>vpls>sap>lag-link-map-profile config>service>vprn>if>sap>lag-link-map-profile config>service>vprn>nw-if>lag-link-map-profile config>service>vprn>sub-if>grp-if>sap>lag-link-mapprofile config>subscr-mgmt>msap-policy>lag-link-map-profile	N
LAG micro BFD	config>lag>bfd	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
LAG per-link hash	config>lag>per-link-hash config>router>if>lag-per-link-hash config>service>epipe>sap>lag-per-link-hash config>service>ies>if>sap>lag-per-link-hash config>service>ies>sub-if>grp-if>sap>lag-per-link-hash config>service>vpls>sap>lag-per-link-hash config>service>vprn>if>sap>lag-per-link-hash config>service>vprn>nw-if>lag-per-link-hash config>service>vprn>sub-if>grp-if>sap>lag-per-link-hash config>subscr-mgmt>sub-profile>egress>lag-per-link-hash	N
LDPv6	config>router>ldp>if-params>if>ipv6 config>router>ldp>if-params>ipv6 config>router>ldp>targ-session>ipv6	N
LLDP	config>port>ethernet>lldp config>system>lldp	N
Load-balancing based on SPI	config>router>if>load-balancing>spi-load-balancing config>service>ies>if>load-balancing>spi-load-balancing config>service>template>vpls-template>load-balancing>spi-load-balancing config>service>vpls>load-balancing>spi-load-balancing config>service>vprn>if>load-balancing>spi-load-balancing config>service>vprn>nw-if>load-balancing>spi-load-balancing	N
Load-balancing based on TEID	config>router>if>load-balancing>teid-load-balancing config>service>ies>if>load-balancing>teid-load-balancing config>service>template>vpls-template>load-balancing>teid-load-balancing config>service>vpls>load-balancing>teid-load-balancing config>service>vprn>if>load-balancing>teid-load-balancing config>service>vprn>nw-if>load-balancing>teid-load-balancing	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Load-balancing based on inner-IP	config>router>if>load-balancing>egr-ip-load-balancing inner-ip config>service>ies>if>load-balancing>egr-ip-load-balancing inner-ip config>service>vprn>if>load-balancing>egr-ip-load-balancing inner-ip config>service>vprn>nw-if>load-balancing>egr-ip-load-balancing inner-ip	N
Load-balancing: LSR options	config>router>if>load-balancing>lsr-load-balancing lbl-ip config>router>if>load-balancing>lsr-load-balancing ip-only config>router>if>load-balancing>lsr-load-balancing eth-encap-ip config>service>vprn>nw-if>load-balancing>lsr-load-balancing lbl-ip config>service>vprn>nw-if>load-balancing>lsr-load-balancing ip-only config>service>vprn>nw-if>load-balancing>lsr-load-balancing eth-encap-ip config>system>load-balancing>lsr-load-balancing lblip config>system>load-balancing>lsr-load-balancing iponly config>system>load-balancing>lsr-load-balancing eth-encap-ip	N
Loop-free alternates	config>router>ip-fast-reroute config>router>ldp>fast-reroute config>router>isis>if>lfa-policy-map config>router>isis>loopfree-alternates config>router>mpls>lsp>igp-shortcut [lfa-protect | lfa-only] config>router>mpls>lsp-template>igp-shortcut [lfa-protect | lfa-only] config>router>ospf>area>if>lfa-policy-map config>router>ospf>loopfree-alternates config>router>ospf3>loopfree-alternates config>router>ospf3>area>if>lfa-policy-map config>service>vprn>isis>if>lfa-policy-map config>service>vprn>ospf>area>if>lfa-policy-map config>service>vprn>ospf3>area>if>lfa-policy-map	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
MAC accounting	config>router>if>enable-mac-accounting config>service>ies>if>enable-mac-accounting config>service>vprn>if>enable-mac-accounting	N
MAC criteria	config>qos>sap-ingress>mac-criteria	N
MAC filters	config>filter>mac-filter	N
Management Ethernet redundancy	config>redundancy>mgmt-ethernet	Y
MDA buffer pool management	config>card>mda>access>ingress>pool config>card>mda>access>egress>pool config>card>mda>network>ingress>pool config>card>mda>network>egress>pool	Y
MLPPP	config>qos>mlppp-profile-egress config>qos>mlppp-profile-ingress config>router>l2tp>group>mlppp config>router>l2tp>group>tunnel>mlppp config>service>vprn>l2tp>group>mlppp config>service>vprn>l2tp>group>tunnel>mlppp config>subscr-mgmt>ppp-policy>mlppp	Partial
MPLS auto-bandwidth	config>router>mpls>auto-bandwidth-multipliers config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth	N
MPLS RSVP P2P LSP secondary path	config>router>mpls>lsp>secondary	N
MPLS RSVP P2MP LSPs	config>router>mpls>lsp <lsp-name> p2mp-lsp config>router>mpls>lsp>p2mp-id config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp-template>p2mp config>router>mpls>p2mp-resignal-timer config>router>mpls>p2mp-s2l-fast-retry	N
MPLS LSP statistics	config>router>ldp>egress-statistics config>router>mpls>ingress-statistics config>router>mpls>lsp>ingress-statistics config>router>mpls>lsp>egress-statistics config>router>mpls>lsp-template>egress-statistics	Partial

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
MPLS strip label	config>router>if>strip-label	Y
MPLS-TP	config>router>mpls>if>mpls-tp-mep config>router>mpls>lsp>protect-tp-path config>router>mpls>lsp>working-tp-path config>router>mpls>mpls-tp config>router>mpls> lsp <lsp-name> mpls-tp	Y
Multi-link protocols	config>port>multilink-bundle	Y
Multicast: MCAC	config>router>mcac config>subscr-mgmt>sub-mcac-policy config>subscr-mgmt>sub-profile>sub-mcac-policy	Y
Multicast: MCS	config>redundancy>multi-chassis>peer>sync>igmp config>redundancy>multi-chassis>peer>sync>mld config>redundancy>multi-chassis>peer>sync>igmp-snooping	Y
Multicast: MLDP	config>router>ldp>mcast-upstream-frr	N
Multicast: L2 snooping	config>service>vpls>pim-snooping config>service>vpls>mesh-sdp>mld-snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>subscr-mgmt>msap-policy>vpls-only-sap-parameters>igmp-snooping BGP-VPLS, BGP-AD, EVPN, and routed VPLS services: config>service>vpls>igmp-snooping config>service>vpls>mesh-sdp>igmp-snooping config>service>vpls>sap>igmp-snooping config>service>vpls>spoke-sdp>igmp-snooping	N
Multicast: MVPN	config>service>vprn>mvpn	N
Multicast: MLDv2 over IPsec	config>service>mld>interface <tunnel-interface-name>	N
Named pools	config>card>named-pool-mode config>card>mda>named-pool-mode config>qos>named-pool-policy	Partial
NAT intra-chassis redundancy	config>isa>nat-group>redundancy config>isa>nat-group>failed-mda-limit	N

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Network ingress and egress policers	config>qos>network>ingress>policer config>qos>network>egress>policer	N
OAM: VPRN ping and trace	oam vprn-ping oam vprn-trace	N
OAM: VCCV ping and trace	oam vccv-ping oam vccv-trace	N
OAM: MAC ping and trace	oam mac-ping oam mac-trace	N
OAM: LSP ping and trace	oam lsp-ping oam lsp-trace	N
OAM: SDP keepalives	config>service>sdp>keep-alive	N
OAM PM	config>oam-pm	N
OAM: MFIB ping	oam>mfib-ping	N
Per-peer queuing	config>service>vprn>bgp-shared-queue config>system>security>per-peer-queuing	N
Persistency	config>system>persistence	N

Table 35 SR OS Features not Supported on VSR (Continued)

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Policy accounting	config>card>fp>policy-accounting config>router>if>ingress>policy-accounting config>router>policy-acct-template config>router>policy-options>policy-statement>default-action>dest-class config>router>policy-options>policy-statement>default-action>source-class config>router>policy-options>policy-statement>entry>action>dest-class config>router>policy-options>policy-statement>entry>action>source-class config>router>static-route-entry>indirect>destination-class config>router>static-route-entry>indirect>source-class config>router>static-route-entry>next-hop>destination-class config>router>static-route-entry>next-hop>source-class config>service>ies>if>ingress>policy-accounting config>service>ies>sub-if>grp-if>ingress>policy-accounting config>service>vprn>if>ingress>policy-accounting config>service>vprn>static-route-entry>indirect>destination-class config>service>vprn>static-route-entry>indirect>source-class config>service>vprn>static-route-entry>next-hop>destination-class config>service>vprn>static-route-entry>nexthop>source-class config>service>vprn>static-route-entry>ipsec-tunnel>destination-class config>service>vprn>static-route-entry>ipsec-tunnel>source-class config>service>vprn>sub-if>grp-if>ingress>policy-accounting	Y
Port buffer pool management	config>port>access>egress>pool config>port>access>ingress>pool config>port>network>egress>pool	N
Port CRC error monitoring	config>port>ethernet>crc-monitor	Y
Port symbol error monitoring	config>port>ethernet>symbol-monitor	Y

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Provider backbone bridging (PBB)	config>service>epipe>pbb config>service>epipe>spoke-sdp>use-sdp-bmac config>service>pbb config>service>sdp>source-bmac-lsb config>service>vpls>b-vpls config>service>vpls>i-vpls config>service>vpls>pbb config>service>vpls>mesh-sdp>fault-propagation-bmac config>service>vpls>sap>fault-propagation-bmac config>service>vpls>spoke-sdp>fault-propagation-bmac config>service>mrp config>service>vpls>isid-policy config>service>vpls>sap>static-isid config>service>vpls>spoke-sdp>static-isid	Partial
PTP	config>router>if>ptp-hw-assist config>service>ies>if>ptp-hw-assist config>service>vprn>if>ptp-hw-assist config>service>vprn>ptp config>system>ptp config>system>sync-if-timing>ptp	Y
PW-Port	config>service>sdp>binding	Y

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
QPPB	config>router>if>ipv6>qos-route-lookup config>router>if>qos-route-lookup config>router>policy-options>policy-statement>entry>action>fc config>router>static-route-entry>indirect>forwarding-class config>router>static-route-entry>next-hop>forwarding-class config>service>ies>if>ipv6>qos-route-lookup config>service>ies>if>qos-route-lookup config>service>ies>sub-if>grp-if>ipv6>qos-route-lookup config>service>ies>sub-if>grp-if>qos-route-lookup config>service>vprn>if>ipv6>qos-route-lookup config>service>vprn>if>qos-route-lookup config>service>vprn>static-route-entry>indirect>forwarding-class config>service>vprn>static-route-entry>next-hop>forwarding-class config>service>vprn>static-route-entry>ipsec-tunnel>forwarding-class config>service>vprn>sub-if>grp-if>ipv6>qos-route-lookup config>service>vprn>sub-if>grp-if>qos-route-lookup	Y
Queue groups	config>card>fp>ingress>access>queue-group config>card>fp>ingress>network>queue-group config>port>ethernet>access>egress>queue-group config>port>ethernet>access>ingress>queue-group config>port>ethernet>network>egress>queue-group config>port>monitor-agg-egress-queue-stats config>qos>queue-group-templates	Partial
SAA	config>saa	N
Schedulers in SLA-profile instance	config>subscr-mgmt>sla-prof>egress>scheduler-policy	N
Selective FIB	config>router>disable-selective-fib config>service>vprn>disable-selective-fib config>system>selective-fib	Y
sFlow	config>port>ethernet>sflow config>sflow	Y

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Shortest path bridging (SPB)	config>service>vpls>sap>spb config>service>vpls>spb config>service>vpls>spbm-control-vpls config>service>vpls>spoke-sdp>spb	N
SONET/SDH	config>port>sonnet-sdh	N
SRRP and MCS	config>redundancy>multi-chassis>peer>sync>srrp config>redundancy>multi-chassis>peer>sync>track-srrp-instances config>router>l2tp>failover>track-srrp config>service>ies>sub-if>grp-if>srrp config>service>vprn>l2tp>failover>track-srrp config>service>vprn>sub-if>grp-if>srrp	N
Subscriber Management: dynamic policers and IP-criteria (RADIUS or Gx)	config>qos>sap-egress>dynamic-policer config>qos>sap-ingress>dynamic-policer	N
Subscriber Management: GTP uplink	--	--
Subscriber Management: queue and policer stat-mode v4-v6	config>subscr-mgmt>sla-prof>ingress>qos>queue>stat-mode config>subscr-mgmt>sla-prof>ingress>qos>policer>stat-mode config>subscr-mgmt>sla-prof>egress>qos>queue>stat-mode config>subscr-mgmt>sla-prof>egress>qos>policer>stat-mode	N
Subscriber Management: subscriber mirroring - including host type criteria	--	N
TMS	config>service>ies>tms-interface config>service>vprn>tms-interface	Y
TDM ports	config>port>tdm	Y
TWAMP	config>test-oam>twamp>server	Y
TWAMP-Light	config>router>twamp-light config>service>vprn>twamp-light config>test-oam>twamp>twamp-light	Y

Table 35 **SR OS Features not Supported on VSR (Continued)**

Feature	Associated CLI (Not Exhaustive)	Commands Removed
Unnumbered interfaces	config>router>if>unnumbered config>service>ies>if>unnumbered config>service>vpls>if>unnumbered config>service>vprn>if>unnumbered	N
uRPF - selective VPRNs	config>router>interface>urpf-selected-vprns config>service>vprn>network>ingress>urpf-check	N
Video ISA	config>isa>video-group config>service>ies>video-interface config>service>vpls>video-interface config>service>vprn>video-interface	Y
VPLS E-Tree	config>service>vpls <service-id> etree	N
VPLS I-PMSI using P2MP LSPs	config>service>vpls>provider-tunnel	Y
Vport scheduling	config>port>ethernet>access>vport	N
VSM	config>vsm	N
Weighted ECMP over RSVP-TE LSPs - VPRN auto-bind	config>service>vprn>auto-bind-tunnel>weighted-ecmp	N
Python for WLAN-GW and NAT groups	config>python>python-policy>nat-group config>python>python-policy>wlan-gw-group	N
WRED	config>qos>slope-policy config>qos>sap-egress>queue>wred-queue	N
WLAN-GW and vRGW: Queue-scale VSR deployment-model	--	--
WLAN-GW and vRGW egress tunnel shaping	config>service>ies>sub-if>grp-if>wlan-gw>egress config>service>vprn>sub-if>grp-if>wlan-gw>egress	N
WLAN-GW and vRGW: pseudowire SAP for L2-AP	--	--
vRGW: PPPoE Client	config>service>vprn>firewall config>service>nat>nat-policy>l2-outside config>subscriber-mgmt>pppoe-client-policy	N
vRGW: IPv6 Firewall	config>service>nat>firewall-policy config>router>firewall config>service>vprn>firewall	N

12 Known Limitations

The following sections describe the known limitations for SR OS Release 19.10.R6.

Note:

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

12.1 Hardware

- The AUX port on the SF/CPM is not supported in software. SR OS does not provide a means of configuring the device.
- The SyncE/IEEE 1588 port on the CPM-a, CCM-e, CPM-2s, 7750 SR-1, and SR-1s 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).
- A SONET/SDH port that is shut down or in internal loopback is incorrectly being allowed as a valid synchronous timing reference. [36448]
- 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]
- 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 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 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]
- Egress FCS-error alarms may in some cases report invalid source card slot numbers for slots that do not have a card equipped. [257024]
- Some 7750 SR-7/12 FP4 configurations could exceed the maximum power capability of the chassis. Refer to the Power Calculator Tool to determine what can be supported in a system. The Power Calculation Tool is available from your Nokia support organization.
- On cpm-x20, holding the reset button for more than four (4) seconds will shut down the CPM and it will not reset. A power cycle of the card is required to get it started again. This can be done either using the **tools**>**perform**>**card** *A* **power-cycle** command, with a physical removal and re-insertion of the CPM, or a power cycle of the entire shelf. [334595]

12.2 Satellites

- If a satellite is being 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 being deconfigured and reconfigured with a new satellite ID, it should first be reset 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]
- Ethernet tunnels are not supported through Ethernet Satellite ports. [282697]
- ETH-CFM MEPs should not be configured with sub-second CCM intervals on a satellite client port using uplink resiliency since the timeout limits for sub-second MEPs are lower than typical resilient switchover times.
- Local-forwarding is not supported on the 7210 SAS-Mxp Ethernet satellite.
- Ethernet Satellite local-forwarding is only supported on the 48x1GE + 4x10GE SAS-S/Sx and 64x10GE + 4x100GE/CFP4 SAS-Sx Ethernet satellites. It is not supported on the 24x1GE + 4x10GE, 64x10GE + 4x100GE/QSFP28 or SAS-Mxp satellite chassis.
- Satellite uplink ports cannot share a 7750 SR IOM with any channelized MDA, including the ASAP and circuit emulation MDAs.
- Configuration of a QoS **slope-policy** is not supported on Ethernet satellite client ports. All Ethernet satellite ports, regardless of the port mode, share a common **slope-policy** associated with the satellite host port. **[NEW]**
- The port **ingress-rate** and **egress-rate** parameters are not supported on Ethernet satellite client ports. A **port-scheduler-policy** should therefore be created to ensure that the 7750 SR OS host is able to shape the traffic according to the egress satellite port type and speed. **[NEW]**
- When mirroring the 7750 SR OS host physical port associated with the uplink to an Ethernet satellite, only satellite control plane traffic is mirrored. To see the traffic associated with client ports, the traffic mirror must be applied to the Ethernet satellite client port. **[NEW]**
- An Ethernet client port cannot be used to receive SyncE timing. Satellite client ports can only transmit SyncE timing. **[NEW]**

- The exponential port dampening feature is not supported by a 7210 SAS operating as an Ethernet satellite. **[NEW]**

12.3 System

- Port-level and SAP-level statistics do not reflect packets processed by the CPM, 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 CPU, 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 system 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 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]

- 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.
- IGMP packets are correctly matched and are dropped or forwarded using a **cpm-filter** entry. However, if logging is enabled, packets are not displayed using the **show**>**filter**>**log** command. [128900]
- 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 interfaces
    - 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.
- Immediately after the **admin reboot** command has been entered, the active CPM may send an event log stating that the standby CPM has failed. This is part of the system shutdown preceding the chassis reboot. [205623]
- 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.
- In an EHS/CRON script, redirecting a command output into a file is not supported. For example, the following command does not work inside an EHS/CRON script.

```
show time > cf3:\output.txt
```

- 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]
- 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
- On a 7950 XRS, when a system boots or a configuration file is executed, if the configuration contains commands under the **card**>**fp** context, but the XMA corresponding to that FP is not provisioned (which may occur if the configuration file is created using the **admin**>**save**>**detail** command in a release prior to 16.0.R1), the first **card**>**fp** command encountered will cause a failure and the processing of the configuration file will terminate. A workaround is to remove all configuration commands under the **card**>**fp** context for FPs that do not have a corresponding XMA provisioned. [283741]
- On 7950 XRS systems, commands can only be issued within the **card**>**fp** context after an XMA corresponding to that FP is provisioned. Conversely, all commands under **card**>**fp** of the corresponding FPs are automatically removed when that hardware is unprovisioned.
- When a *system-name* is configured with **configure system name** *system-name* with a string longer than 32 characters, the following should be considered. [288438]
    - For a DHCPv6 relay, a group interface can be configured to insert the *system-name* as part of the Interface-ID Option 18 (**dhcp6 option interface-id ascii-tuple**).
    - An SR OS DHCPv6 server drops Relay-Forward messages when the Interface-ID is longer than 64 characters.
- When saving the configuration at the start of a software upgrade, the **detail** option of **admin save** (for example, **admin save detail**) should not be used on the VSR platform. Upgrades with **admin save detail** are not supported on VSR.
- On the me2-100gb-ms-qsfp28 MDA, the **max-peer** attribute under the MACsec sub-port is forced to a value of one (1). Only one MACsec peer is allowed per port. [318282]
- On the me2-100gb-ms-qsfp28 MDA, only one MACsec sub-port can be configured per connector breakout. On a 4-port breakout, adding a MACsec sub-port on a subsequent port under the same connector breakout will fail. [318283]

12.4 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.

12.5 Diameter

- The following Diameter protocol restrictions apply:
  - Accounting (RFC 6733) using Diameter is not supported
  - Accounting-Request (ACR), Accounting-Answer (ACA), Session-Termination-Requests (STR), and Session-Termination-Answer (STA) messages are not supported
  - SCTP and IPsec as transport protocols are not supported. TCP is supported.
  - Diameter Proxy in Multi-Chassis Environment is supported only on legacy Diameter base
- A mix of single homed (non-synchronized) and dual-homed (synchronized) Diameter-enabled subscribers is not supported in Diameter multi-chassis redundancy. All Diameter-enabled subscribers must be dual-homed when Diameter multi-chassis redundancy is enabled for a specific Diameter node.

12.6 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]

12.7 Classic CLI

- Non-printable, 7-bit ASCII characters or strings only containing spaces are not allowed inside the various description fields. [93998, 284419]
- 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 failover occurs (the candidate will be empty).

12.8 Model-driven Interfaces

12.8.1 Common Limitations

This section lists limitations that apply to all MD interfaces.

- Configuration via model-driven (MD) interfaces (the MD-CLI, NETCONF, gRPC) is not supported for the entire SR OS configuration data model. The existence of configuration elements in the Nokia SR OS MD interfaces in a given release does not mean all those elements are fully functional and available for transactional configuration use in all scenarios. Nokia recommends avoiding any use of unsupported configuration elements in deployments. See Unsupported Configuration in MD Interfaces for information about configuration coverage in MD Interfaces.
- State information modeled in Nokia SR OS YANG models, and available via Telemetry (subscriptions) and the NETCONF and gNMI (gRPC) model-driven (MD) interfaces, is supported for a subset of SR OS state. The existence of state elements in the Nokia SR OS MD interfaces in a specific release does not mean all those elements are fully functional and available in production networks. Nokia recommends avoiding any use of unsupported state elements in deployments, and avoiding state retrieval using paths near the top of the state tree that include large sections of the state data model. Contact a Nokia representative for information on the availability of specific state information in a specific release.
- The following items are not supported when **configure system management-interface configuration-mode** is set to **model-driven** or **mixed**:
  - Dynamic Data Services
  - Configuration control from the Nuage VSD via XMPP
  - SONET/SDH, ATM, Frame Relay, TDM, etc. interfaces. Only directly connected physical Ethernet-based interfaces and Satellite-based Ethernet interfaces are supported.
  - Loose references to IDs for several types of configuration objects. For more information, see "Loose References to IDs" in the *System Management Guide*.
  - Full management of configuration via SNMP (including via the Nokia NSP products, for example, using NFM-P with SNMP-based configuration control) is not supported in either model-driven or mixed configuration modes. SNMP writes (for configuration or for actions such as node reboot) are not supported in mixed or model-driven configuration-mode. SNMP reads are supported in all configuration modes.

- The following items are not supported when the primary CLI engine set by **configure system management-interface cli cli-engine** is set to **md-cli**:
    - .if/.set commands within an EHS script
- The following items apply to the use of the Nokia SR OS YANG modules/ submodules and data model (XML namespace urn:nokia.com:sros:ns:yang:sr:conf-* and state-*)
    - The Nokia SR OS YANG configuration data model is composed of a single module and multiple submodules. Some YANG tools may complain about circular dependencies in the submodules. Pyang complains about circular dependencies but does complete the processing to build complete tree or jstree output. If circular dependencies are preventing any necessary tools from correctly processing the YANG, then use the alternative packaging of the Nokia configuration model available in the single file called nokia-conf-combined.yang. The same limitation and solution applies to the Nokia SR OS YANG state data model.
    - The Nokia SR OS YANG modules are expected to be updated in subsequent releases without adhering to all of the module update rules specified in RFC 6020 Section 10. Some changes will not be backward 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 additional containers may be added, or leafs moved between containers (mostly to clarify, improve consistency, or fix errors)
        - Some objects may change from leafs to containers or 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 “choice” constructs, “must” conditions, “mandatory” statements etc.)
        - Some strings may change to leafrefs. This is primarily for references to objects that were not previously part of the subset of configuration that was supported in the Nokia SR OS YANG modules but are now supported.
        - Larger submodules may be sub-divided into several smaller submodules for better modularity and clarity.
        - Some default values may change, be removed, or be added.
        - The type of some leafs may change from string to leafref with require-instance false (a YANG 1.1 construct).

- The type of some leafs may change from string to a union of leafrefs (a YANG 1.1 construct).
- YANG 'if-feature' statements may be added to describe applicability of certain schema nodes.
- Some integer types may be expanded (for example, from uint32 to int64) in order to accommodate larger values.
- Some boolean leafs may be changed to enums to accommodate additional configuration values besides true and false.
- Some enum values may be removed (for example, for hardware items that are no longer supported).

- The existence of configuration elements in the OpenConfig model in a specific release does not mean all those elements are fully functional and available for transactional configuration in all scenarios. Deviation files contain additional information for the supported models. Contact your Nokia representative to determine if your configuration needs are available via OpenConfig models in a specific release.
- In NETCONF and gNMI (gRPC) model-driven interfaces, AAA command authorization does not operate for YANG state information. Access to state paths are not controlled and all state paths and leafs are available to any authenticated user.
- Operations such as **load**, **commit**, and **admin save** (and their equivalents in NETCONF or gNMI) in model-driven interfaces exhibit longer durations that may be noticeable in highly-scaled configurations as compared to their equivalent in classic-configuration mode. This is mainly related to the additional processing required to achieve the advantages of transactional configuration.
- The YANG “must” statements introduced in Release 19.10.R1 and then removed in Release 19.10.R3 may cause problems integrating with third-party NETCONF clients or controllers. It is recommended to use other releases to integrate with third-party NETCONF clients or controllers.

12.8.1.1 System

- Model-driven or mixed-management interface configuration mode is not supported with the following non-Ethernet MDAs: [305353]
  - m1-choc3-ces-sfp
  - m1-choc12-as-sfp
  - m1-choc12-ces-sfp
  - m2-oc192-xp-xfp
  - m4-chds3-as

  - m4-choc3-as-sfp
  - m4-choc3-ces-sfp
  - m4-oc48-sfp-b
  - m12-chds3-as
  - m16-oc12/3-sfp-b
  - vsm-cca-xp
- A telemetry ON_CHANGE subscription for configuration events and the mgmt_core application's mdConfigChange log event incorrectly report inherited configuration elements instead of the **apply-groups** element when applying a configuration group.
- Lawful Intercept (LI)
  - Configuration groups
  - When **li-separate** mode is enabled, executing MD-CLI **rollback** and **load** commands in the main configuration region are denied with a "Permission denied" error message.
- A router booted with empty configuration, without a card and port configuration applied, can afterwards fail to commit **router interface ip-filter** configuration that is executed via **load full-replace**. This will result in an error message: "The port has not been configured”. [335014]

12.8.1.2 Routing

- When ECMP is enabled and one route is learned using BGP and the second route is learned using BGP VPN, the protocol of the route is mentioned twice as BGP when reading the route-table state using MD interfaces. [305866]
- In mixed-management interface configuration mode, non-default BFD configuration parameters set with a model-driven interface are removed from the configuration if BFD is configured as **shutdown** or **disabled**. [306702]
- When changing a static security association (=**static-sa**) used for OSPFv3 authentication, the rollover timeout (**key-rollover-interval**) must be taken into consideration. During security association rollover, the previous old **static-sa** can still be valid for up to 300 seconds (the default timer value is 10 seconds) and the previous **static-sa** is still used. During the rollover time deleting the old **static-sa** will not be possible. If, however, during rollover time there is an attempt to delete this old **static-sa**, an error message will be spawned [MINOR: IPSEC_CPM #1004: configure ipsec static-sa "OLD_POLICY_FOO" - Static Security Association cannot be deleted while it is in use by a policy]. The old **static-sa** can only be deleted once the rollover-process interval timer expires. The consequence of this technology property is that changing a configured OSPFv3 security association needs to be split into two parts: [311398]

1. Firstly, the new **static-sa** has to be configured and activated for OSPFv3.
2. Secondly, after the **key-rollover-interval** has expired, the old **static-sa** can be deleted.

12.8.1.3 MPLS

- The following state leafs should not be used. They were introduced in 19.10.R1 and were removed from the YANG models (not deprecated or obsoleted) in 19.10.R3:
  - **state router mpls ipv4-oper-state**
  - **state router mpls ipv6-oper-state**
  - **state router mpls interface ipv4-oper-state**
  - **state router mpls interface ipv6-oper-state**

  Instead, use the following state leafs:
  - **state router mpls ipv4 oper-state**
  - **state router mpls ipv6 oper-state**
  - **state router mpls interface ipv4 oper-state**
  - **state router mpls interface ipv6 oper-state**
- In MPLS, *lsp-name* (for example, in **configure router mpls lsp** *lsp-name*) and *lsp-template-name* (for example, in **configure router mpls lsp-template** *lsp-template-name*) must be unique when configured in a model-driven interface. The system checks the configuration and validates against other configured names, but it cannot check against dynamically-created LSP names during validation. The existence of an *lsp-name* (or *lsp-template-name*) that matches with a dynamically-created LSP name is only detected during a commit and causes the commit to fail. [310261]

12.8.1.4 Services

- All services that reference a port or interface defined using the OpenConfig (OC) model must be deleted and committed in a separate transaction before the OC configuration for the associated OC interfaces can be deleted.
- When a transaction results in the deletion of a **pw-template-binding** within a VPLS or Epipe service, the commit fails if the **pw-template-binding** is in use. The user must remove all active PW bindings that are using this **pw-template-binding** for the commit to succeed. [297555, 319007]

- A single transaction that involves any change in operational route distinguisher value that affects one or more service or BGP instance is not supported. The desired configuration change must be performed using multiple individual transaction commits. [300174]
- An L2TPv3 cookie value cannot be added under a spoke-SDP in the same transaction while the underlying SDP is being changed to L2TPv3. This must be performed as multiple transactions. [305672]
- Changing the **bgp-auto-rd-range** may require deleting the existing **bgp-auto-rd-range** in one transaction (which may also require deletion of all service route-distinguishers using the automatic range), and then adding the new **bgp-auto-rd-range** in a second transaction. [311499]
- Changes in a **pw-template** requires running the **tools perform service eval-pw-template** *policy-id* **allow-service-impact** command so the changes are effective in the PWs of a service. This is the same behavior as in classic configuration mode.

12.8.1.5 Subscriber Management

- In the MD-CLI, the **tools perform subscriber-mgmt coa** command has the following limitations when specifying attributes on the command line with **attr**: [295198]
  - Attributes with their values must be enclosed in double quotes. For example, **tools perform subscriber-mgmt coa nas-port-id "1/1/c1/6:1000.1" delegated-ipv6-pfx 2001:db8:0:55::/64 attr ["6527,13=sla-profile-1" "6527,126=e:q:1:pir=20000,cir=5000"] debug**
  - Double quotes cannot be used in the attribute value or to enforce a string format for an attribute. A workaround is to specify the attributes in a text file instead.
- ESM is not supported on services using an auto-generated service ID.
- ESM hosts do not support filters with an auto-generated filter ID.
- When modifying a configuration that is in use by active subscriber management hosts, it is possible that the transaction will remove these hosts, the transaction will be blocked, or a combination of both. For example:
  - Changing the **group-interface** *group-interface-name* **ipv6 router-solicit inactivity-timer** will result in a failed transaction if any IPv6 host is active.
  - Changing the **subscriber-mgmt ipoe-session-policy** *name* **circuit-id-from-auth** flag will result in removing of all IPoE sessions on group interfaces using this policy.
- The **apply-groups** configuration element applied to configuration groups should not be configured on the following branches: **[NEW]**

- **configure service vpls** *service-name* **sap** *sap-id* **sub-sla-mgmt**
- **configure service ies** *service-name* **subscriber-interface** *interface-name* **group-interface** *interface-name* **sap** *sap-id* **sub-sla-mgmt**
- **configure service vprn** *service-name* **subscriber-interface** *interface-name* **group-interface** *interface-name* **sap** *sap-id* **sub-sla-mgmt**

12.8.1.6 NAT

- In the classic CLI, NAT static port forwarding (SPF) can be done using configuration commands (saved in the configuration file) or using tools commands (exists as state, only saved to compact flash if **nat-port-forwarding** persistence is enabled). In MD Interfaces, managing NAT SPF using configuration is not supported. NAT SPF can only be done in the MD-CLI using the **tools perform nat port-forwarding-action** command.

12.8.1.7 WLAN-GW

- For a **radius isa-policy** *policy-name* **servers source-address-range**, there is a maximum of 14 (maximal members in a system) addresses following the starting address where the IP address is allocated. This must be considered when doing a transaction with model-driven interfaces where the **source-address-range** potentially overlaps with another IP address. The system cannot always detect overlap before executing the transaction, which may cause a transaction failure and rollback. [320404]
- When modifying a configuration that is in use by active WLAN-GW UEs, it is possible that either the transaction will remove these UEs, the transaction will be blocked, or a combination of both. See the Subscriber Management section of Known Limitations for examples.

12.8.2 MD-CLI

This section lists limitations that apply only to the MD-CLI.

- The MD-CLI is a model-driven (MD) interface and is subject to the general limitations of MD interfaces. See the Model-driven Interfaces section of Known Limitations for more information.

- The maximum MD-CLI input line length is too short for OpenConfig leaf-lists that use the maximum string length for all elements. This configuration is unlikely to occur in a production network. [309005]
- The MD-CLI **?** help displays the minimum and maximum values defined in the YANG model range statement, which may be different than the values that are applicable and accepted by configuration commands.

12.8.3 gRPC

- The gRPC server supports JSON and BYTES encoding. All other encodings defined in gnmi.proto (version 0.7.0) are not supported.
- The Alias concept (as defined by gnmi.proto version 0.7.0) is not supported.

12.8.3.1 gNMI

- The gRPC gNMI Interface is a model-driven (MD) interface and is subject to the general limitations of MD interfaces. See the Model-driven Interfaces section of Known Limitations for more information.

12.8.4 Unsupported Configuration in MD Interfaces

All configuration elements (for example, under **configure** in the classic CLI) that are supported in the classic CLI are also supported in MD interfaces except the items listed below.

The following sub-sections list configuration items in the classic CLI that are not supported in MD interfaces.

12.8.4.1 System

- Bluetooth
- Non-Ethernet interfaces and related functionality:
  - SONET/SDH interfaces
  - OTU configuration
  - Channelized interfaces and TDM configuration

  - Bundles (for example, MLPPP)
  - APS
  - Cpipe services
  - ATM (interfaces, Apipes, ATM QoS, **connection-profile**, etc.)
  - Frame Relay (interfaces, Fpipes, FR QoS etc.)
- PTP (IEEE 1588) in **model-driven** configuration mode and via MD interfaces (PTP configuration is available in **mixed** configuration mode using classic interfaces only)
- TDM satellites
- Satellite **local-forward**
- sFlow
- XMPP for Nuage VSD integration
- **boot-bad-exec**

12.8.4.2 OAM

- Ethernet CFM MIP creation
- SAA: All test types except **icmp-ping** (which is supported in Release 16.0)
- SAA: **trap-gen**
- ETH-CFM QinQ tunnel MEPs (Service and SAP/SDP **tunnel-fault**)
- **test-oam build-packet** and **ldp-treetrace**
- Ethernet ELMI

12.8.4.3 Routing

- Router **policy-acct-template**
- Router **service-prefix**
- CPM router instances (user-named instances in **configure router**)

12.8.4.4 Multicast

- **router tunnel-interface** (for P2MP LSPs)

12.8.4.5 MPLS

- MPLS-TP
- GMPLS

12.8.4.6 Services General

- XMPP-based Data Center Gateway integration (Static-Dynamic and Fully-Dynamic models) with the Nuage Virtual Services Directory (VSD). When MD interfaces are used, NETCONF should be used to integrate the system with VSD.

12.8.4.7 L2 Services

- B-VPLS MRP, and MMRP
- SPB
- **eth-ring** (ITU-T G.8032)
- Mesh-SDP/SAP/Spoke-SDP **fault-propagation-bmac**
- Epipe/VPLS load balancing
- **configure service system vxlan**
- **configure service vpls mac-notification**
- IGMP snooping in a VPLS VXLAN service
- MLD snooping in a VPLS VXLAN service
- Ipipes
- **eth-tunnel** (ITU-T G.8031)
- **configure service epipe sap ethernet llf** (Link Loss Forwarding)
- VPLS and VPLS SAP templates
- VPLS **mc-endpoint**
- VPLS interface
- **vpls sap force-l2pt-boundary**
- **vpls sap dest-mac-rewrite**
- PW Routing
- MVRP (including **vpls-group**)

12.8.4.8 L3 Services

- IES IPCP
- **configure service ies|vprn interface ipv6 dhcp6-server prefix-delegation**

12.8.4.9 QoS

- Card
  - FP ingress access queue group policer stat-mode no-stats
  - FP ingress network queue group policer stat-mode no-stats
- Port
  - Access ingress pool
  - Egress queue group policers
  - Ethernet egress access queue group **agg-rate**
  - **monitor-agg-egress-queue-stats**
  - **eth-bn-egress-rate-changes**
- QoS Policies
  - Scope exclusive in SAP ingress, SAP egress and network QoS polices
  - Egress queue group template policers
  - Network policy ingress FC
  - Network policy egress IP/IPv6 entry action
  - Shared queue
  - Named pool policies
- Services
  - Applied SAP ingress QoS
    - with MAC criteria statements applied to routed VPLS SAPs
  - Applied SAP egress QoS with the applied QoS policy configured with port egress queue group redirection and applied under
    - Epipe PW-SAPs
    - IES interface PW-SAPs
    - VPRN interface PW-SAPs
  - Scheduler policy applied to
    - IES interface SAP ingress
  - Policer control policy applied to
    - IES interface SAP egress

    - Ingress and egress QoS applied to
        - Epipe SAP templates
        - VPLS SAP templates
        - PW-templates
    - **match-qinq-dot1q** under
        - Epipe ingress SAPs
        - IES interface ingress SAPs
    - **qinq-mark-top-only** under
        - IES interface egress SAPs
    - PBB VPLS B-SAP per-ISID shaping
    - VPRN Network interface QoS
    - PW-SAP secondary shapers on an IOM4-e-HS and HS-MDAv2
    - Ingress and egress PW QoS for Epipe spoke SDPs, VPLS spoke and mesh SDPs, IES interface spoke SDPs and VPRN interface spoke SDPs.

12.8.4.10 IPsec

- Secured interface

12.8.4.11 Subscriber Management

- SHCV configuration without policies
- RADIUS server configuration without **radius-server-policy**
- Dynamic Data Services
- RADIUS Python scripts without Python policies
- **accept-authorization-change** in an **authentication-policy**
- VPLS SAP templates for use in WLAN-GW Layer-2 wholesale
- NGE Support for WLAN-GW L2oMPLSoGRE

12.8.4.12 VSR

- VSR Flex PW-port support (including **pw-port-list**)
- Filter **ip-exception** (for NGE, on SAR-Hm and VSR only)

12.8.4.13 Additional Items

- Application Assurance
- Video

12.9 Mixed Management Interface Configuration Mode

- When operating in **mixed** configuration mode, configuration in classic CLI is subject to the same availability and coverage constraints as model-driven interfaces. Only configuration elements that are fully available in model-driven interfaces can be edited in the classic CLI. The classic CLI cannot be used to configure elements that are not supported yet in MD interfaces. See Unsupported Configuration in MD Interfaces for the availability of specific configuration elements in this release.
- Saved configuration files generated in **classic** configuration mode may be incompatible with **mixed** configuration mode. Nokia recommends running **admin save** to save the configuration file after switching to **mixed** configuration mode.
- The following items are not supported when **configure system management-interface configuration-mode** is set to **mixed**:
  - OpenConfig YANG models
  - Configuration groups
  - Named route policy entries

12.10 Telemetry

- State information available via Telemetry has some limitations. See the Model-driven Interfaces section of Known Limitations for more information.
- Under overload conditions (that is, the output queue with Notification messages internally defined threshold), some ON_CHANGE notifications may be dropped. This typically happens if the ON_CHANGE notification is referring to an object which might have been deleted from the configuration. [329934]
- The POLL subscription mode (as defined by gnmi.proto version 0.7.0) is not supported.

12.11 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.
- When an FP3-based line card is provisioned in a 7750 SR-12e equipped with SFM6-12e and only FP4-based line cards, and the **configure mcast-management chassis-level per-mcast-plane-capacity total-capacity** is set to **dynamic** (the default), the IMPM path/plane capacity will reduce from 19 G to 8.25 G. Conversely, when the last FP3-based line card is unprovisioned leaving only FP4-based line cards in the system, the IMPM path/plane capacity will increase from 8.25 G to 19 G. The chassis multicast capacities can be displayed using the **show mcast-management chassis** command.
- When an FP3-based XCM is provisioned in a 7950 XRS-20 or 7950 XRS-20e equipped with only FP4-based XCMs and XMAs, and the **configure mcast-management chassis-level per-mcast-plane-capacity total-capacity** is set to **dynamic** (the default), the IMPM path/plane capacity will reduce from 15 G (in a 7950 XRS-20) or 19 G (in a 7950 XRS-20e) to 8.25 G. Conversely, when the last FP3-based XCM is unprovisioned leaving only FP4-based XCMs and XMAs in the system, the IMPM path/plane capacity will increase from 8.25 G to 15 G (in a 7950 XRS-20) or 19 G (in a 7950 XRS-20e). The chassis multicast capacities can be displayed using the **show mcast-management chassis** command.

12.12 SONET/SDH

- 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]
- 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, 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.

12.13 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]

12.14 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]

12.15 PPP

- PPP is not preventing IPCP negotiation with a non-matching IP subnet address. [24475]

12.16 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 types. [143700]
- 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.

12.17 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 CPM High-Availability 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 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]

12.18 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 BFD packets use hashing independent of that configured for the data traffic. When **per-fp-egr-queuing** is enabled, the BFD packets may egress LAG over a different port than used by the SAP's data traffic. For those BFD packets, internal system queues, instead of the SAP's queues are used, and BFD packets are not accounted for in the SAP queues.
- Pulling out the active CPM 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 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.
- When traffic is sent over PW-port SAPs that are configured on a LAG with the **per-link-hash weighted** option enabled, the traffic is dropped. The CLI incorrectly allows this configuration combination.
- For Epipe using **per-service-hashing** load balancing, the **tools dump map-to-phy-port lag** command does not return the expected egress physical link in case the ingress traffic is subject to QoS policing. [308656]

- VSR always includes L4-port information and the system-IP in its hash calculations for ECMP and LAG load-balancing purposes, irrespective of configuration.

12.19 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]

12.20 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 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 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's APS controller (as displayed in CLI). [36537]

12.21 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]

12.22 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]
- SNMP traps are not forwarded when overwriting or modifying an existing trap-target in both the base and VPRN contexts. [177129]

12.23 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 CPM 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:
  - 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.
- BGP best path selection always prefers an IPv6 next-hop over an IPv4 next-hop and indicates a TieBreakReason of NHType. ECMP across a mix of IPv4 and IPv6 BGP next-hops requires configuring **ignore-nh-metric** under best-path-selection.
- The GRE termination subnet is not supported with the following interfaces:
  - an unnumbered network IP interface
  - an IES interface
  - a VPRN interface
  - a CSC VPRN interface
- Non-segmented MLDP intra-AS option B and its ABR MoFRR are not supported in Release 16.0.
- The SR OS implementation does not support redirection of IPv6 flows, redirection towards an IPv6 address, or any copy (C=1) functionality. There is also no support for interpreting the next-hop address in the MP_REACH_NLRI attribute of a received FlowSpec route as a redirection address.
- BGP FlowSpec Support for Redirect-to-LSP Action has the following limitations:
  - Only a single redirect-to-LSP extended BGP attribute per FlowSpec NLRI is supported and consequently the "S-ID" (=sequencing-ID) is ignored.
  - FlowSpec redirect-to-LSP within a VPRN is not supported. This type of implementation redirection within a VPRN risks of escaping of data outside of the VPRN. The functionality is blocked, and a log message iscreated.
  - Only redirection ID-Type "0" (localised ID) is supported. Other redirection IDs result into an event log warning.
  - FlowSpec redirect-to-LSP with the 'C' (=copy) bit set is not supported. A warning is displayed in the log to identify the unsupported requested action.

- If a TCP connection is already established before an SR OS node is upgraded to Release 16.0.R4 or higher, and TCP-AO is enabled for that pre-existing TCP connection, the TCP connection and associated peering will be reset due to a mismatch in the parameters maintained for a TCP session prior to Release 16.0.R4.
- The default value for **ebgp-default-reject-policy** depends on the configuration management interface, either classic (the classic CLI and SNMP) or model-driven (the MD-CLI, NETCONF, or gRPC) that created the BGP instance. A group or neighbor default value can be configured over different management interfaces, but is not recommended because the value is inherited from the BGP instance. This can cause a different operational BGP behavior after a node reboot. [316955]
- The following limitations apply to gRPC-based RIB/FIB API:
  - If a gRPC client using the RIB-API service omits any parameter that is considered mandatory by the server-side of the RIB API service, the router assumes that the intended value for the parameter is zero (0). This may cause an error if the zero value is invalid or unavailable.
  - If a route, tunnel or MPLS label entry is modified, and it is covered by an ON_CHANGE subscription to a state path enabled by **fib-telemetry**, the update will replay the current values of the entire entry (except for statistics), including values did not change from the last update. It is up to the client to compare the update to the previous one received if it needs to know the exact properties that changed.
  - A gNMI telemetry subscription cannot mix state paths enabled by **fib-telemetry** with other state paths outside of this area.
  - gNMI telemetry subscriptions to list keys of state paths enabled by **fib-telemetry** are not supported.
  - A gNMI telemetry subscription enabled by **fib-telemetry** may provide incorrect details for routes or tunnels programmed through means other than the RIB API.
  - Shutting down gRPC on the router when it is in the process of sending or receiving a protobuf messages can cause the server to end up in a bad state. It is only recoverable by a reboot of the node or a CPM switchover. [324579]
- The following limitations apply to RFC 6549 OSPFv2 Multi-Instance Extensions:
  - OSPFv2 multi-instance extensions support supports using the same GRT interface in different OSPFv2 instances, similar to the multi-instance support for OSPFv3.
  - Multi-instance OSPFv2 is only supported for BASE routing, and not for VPRN services. Support for multiple OSPF instances in VPRN is not supported.

  - Support for OSPFv2 multi-instance Extension (RFC 6549, OSPFv2 Multi-Instance Extensions) does not change this behavior.
  - An OSPF router in instance 0 with multi-instance enabled forms an OSPF adjacency with another router in instance 0 without multi-instance enabled.
  - Support for multi-instance (RFC 6549) is backward compatible.
  - If there are multiple instances (1, 2, and 3) on a first router with multi-instance enabled, and only 1 instance on another router without multi-instance enabled, causes adjacency errors on only one instance. This will be the first instance found in the tree and is not necessary the lowest instance ID.
- Multicast or BIER together with forwarding adjacency in single topology IS-IS (MT0) is not supported. Multicast must be routed using multicast multi-topology IS-IS topology (MT3). [325576]

12.24 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]

12.25 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]

12.26 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
- Received and self-generated IPv6 packets are dropped when they should be forwarded from an interface that does not meet the minimum required Layer-3 MTU of 1280 bytes for an IPv6 interface. [115437]

12.27 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]
- When configuring **dhcp6-relay** and specifying the client global IPv6 interface address as **source-address**, the Relay Reply message from a DHCPv6 server can be dropped with the reason "Relay Reply Msg on Client Itf". Using a client global IPv6 interface address as **source-address** is not supported in this type of setup, but an interface loopback IPv6 address can be used. [292642, 313238]

12.28 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]

12.29 IS-IS

- IS-IS is not supported on IES and VPRN interfaces with ATM PVC SAPs.
- 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 switching a keychain entry, the database has LSP entries with the old entry (used as *key-id* in the authentication TLV). If an adjacency then comes up, it receives these old entries with a *key-id* that may no longer be valid if the tolerance is exceeded. This prevents the adjacency from coming up. When all LSPs are refreshed, they all have the new *key-id* and the adjacency will form. Using a tolerance larger than the **lsp-refresh-interval** will prevent this occurrence. When using **auth-keychain** in IS-IS, Nokia recommends setting the keychain's entries' tolerance time to at least the IS-IS **lsp-refresh-interval**. [248372]
- 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]
- For valid BIER transport between two ABRs (dual-level capable routers), the Bit-Forwarding Routers (BFRs) between them should have their BFR-prefixes in Level-1 (regardless it is present in Level-2 or not). [305679]

12.30 OSPF

- The system may refresh self-originated LSA shortly after completing a CPM 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 a CPM 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]

12.31 OSPF PE-CE

- OSPF traffic engineering is not supported in VPRN instances.

12.32 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 CPM 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 RFC 3107-labeled 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 correct way to trigger GR is to use the **clear router bgp neighbor** *ip-address* 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 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.
- 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.

- The **bgp-ad no shutdown** and **bgp-vpls no shutdown** commands are incorrectly restricted if an SDP-binding has **force-vlan-vc-forwarding** enabled. However, using **vc-type vlan** on **pw-template** for **bgp-ad** and **force-vlan-vc-forwarding** on manual SDP-bindings in the same VPLS service is supported. [281403]
- When BGP on the router advertises FlowSpec routes to EBGP peers, non-transitive extended communities are not stripped from the advertised routes.
- The resolution of a /32 label-ipv4 route with a prefix-SID attribute does not create a tunnel in the Segment Routing (SR) database; it only creates a label swap entry when the route is re-advertised with a new next-hop. This means that the first SID in any SID-list of an SR policy or of a SR-TE LSP should not be based on a BGP prefix SID. If the first SID is based on a BGP prefix SID, the SID-list may appear to be valid but the datapath will not be programmed correctly. However, it is fine to use a BGP prefix SID as any non-first SID in any SR policy. [377255]

12.33 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 ETH-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
  - **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 **config**>**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.
- When two BGP instances are enabled on the same VPLS service, the following features are not supported:
    - SDP-bindings
    - M-VPLS, I-VPLS, B-VPLS or E-Tree 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 multihoming 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 **config**>**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.

- A BGP route-reflector may reflect a wrong value in the EVPN routes Label field, if the routes are received without BGP encapsulation extended community, and the **def-recv-evpn-encap** command is configured with a different encapsulation than the service that advertised the route. For example, assuming a route-reflector is configured with **def-recv-evpn-encap mpls**, an EVPN-VXLAN route received without BGP encapsulation extended community will be reflected with the low-order four bits set to zero (as in a regular MPLS label field). The remote PE will therefore decode a different VXLAN VNI value than the originating PE advertised. [283005]
- The **bgp-evpn**>**ip-route-advertisement** command is not supported along with the following combinations in a R-VPLS service:
  - Two static VXLAN instances in the R-VPLS service
  - One static VXLAN instance and one BGP-EVPN VXLAN instance (note that the command is not blocked in this case.)
  - One static VXLAN instance and one BGP-EVPN MPLS instance (note that the command is not blocked in this case.)
  - Two BGP-EVPN VXLAN instances
  - One BGP-EVPN VXLAN instance and one BGP-EVPN MPLS instance
- R-VPLS interfaces configured as **evpn-tunnel** are not supported when the associated VPLS uses BGP instance 2. The system does not restrict this configuration.

12.34 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.

12.35 PCEP

- The LSP name string of a PCC-initiated LSP or of a PCE-initiated LSP must not contain a single-column character ":" or a double-column character "::". SR OS uses "::" as the delimiter for concatenating the LSP name and path name of RSVP-TE LSP or SR-TE LSP in the LSP database. [315230]

12.36 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 an SR-TE LSP with a path computed with the hop-to-label method. This is not an issue for a SR-TE LSP which path is computed by the local CSPF or by PCE.
- When a SR-TE LSP path uses an OSPF P2P link for which is not enabled in a RSVP context, the remote address of the link shows as *0.0.0.0* in the output of the path show command and lsp-ping and lsp-trace probes will fail. The workaround is to enable RSVP on the OSPF interface.
- The IPv6 SR-TE LSP does not support the following features:
  - Local CSPF and PCE path computation methods. Only the hop-to-label translation method is supported.
  - LSP template
  - Seamless BFD

- An IPv4 or IPv6 SR-TE one-hop auto-LSP, as well as a configured IPv4 or IPv6 SR-TE LSP with a first hop matching a local adjacency, remains down when **path-computation** is set to the **local-cspf** value and the neighbor node does not enable Segment Routing.

12.37 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]
- 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 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.
- 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.

- LSP BFD over tunnels to non-host prefixes is not supported. Only BGP IPv4 prefixes in GRT can make use of such tunnels as part of BGP next-hop resolution. This capability is not supported for BGP IPv6 in GRT, VPN-v4, VPN-v6 label-IPv4, and label-IPv6 routes.
- If a configuration with a reserved label block is saved via **admin rollback save**, and next a reserved label block range is changed by the user and the system dynamically allocates a label value that was previously reserved for the reserved label block, then a subsequent **admin rollback revert latest-rb** will fail. [288535]
- Values for mplsInterfaceTotalBandwidth and mplsInterfaceAvailableBandwidth are defined in the standard MPLS-LSR-MIB as Integer32 and are in Kb/s. If the total bandwidth or available bandwidth on the interface exceeds 2147483647 Kb/s, these objects will show incorrect values. The system stores the correct values in vRtrRsvpIfBandwidth and vRtrRsvpIfReservedBandwidth, both of which are of type Unsigned32. These are available via SNMP, **show** commands, and telemetry paths. [317961]
- If MPLS forwarding policy egress statistics collection is enabled on more constructs than the system's limit, it is not possible to control which constructs will receive the statistical indices and which constructs will not because the allocation of statistical indices after a reboot is not deterministic.

12.38 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.

12.39 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 CPM 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. 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 High Availability switchover occurs, the running **ldp-sync** timer could be incorrectly decremented, inducing an early activation of the static-route. [224939]

12.40 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, and BGP-VPWS) 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 FEC bonding capabilities are not supported with LDP IPv6:
  - resolution of IPv6 FEC over an RSVP IPv4 LSP

12.41 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.

12.42 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

12.43 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]
- PIM-snooping **mode snooping** is not supported in a dual-homed EVPN-VXLAN scenario where MCS is used to synchronize the multicast states across the two Ethernet Segment (ES) peers. Only **mode proxy** is supported in this scenario. [320969]
- The **apply-bgp-nh-override** CLI option in the **config**>**service**>**vprn**>**pim** context is only available in MVPN. This command is used to force the RPF to be checked against the IPv4 VPN AF next-hop.

12.44 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.
- 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]

12.45 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 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 **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.
- Assigning a multi-service site to a LAG configured with **adapt-qos distribute** is not supported when the LAG ports span multiple FPs on a 7950 XRS XCM or a 7750 SR-14s XCM-s.
- When policers are parented to a root or intermediate arbiter in a SAP ingress or egress QoS policy which is applied to a LAG SAP, where the LAG is configured with **adapt-qos distribute** and the LAG ports span FPs on a 7950 XRS XCM, a 7750 SR-14s XCM-s, or a multi-FP XMA, a set of policers and (root and intermediate) arbiters will be instantiated per FP.

12.46 QoS (VSR)

- Orphan queues (without a parent or port-parent) are serviced by a default port scheduler. This default port scheduler associates each orphaned queue with a level equal to its queue id.

- For egress traffic, VSR always uses the in-profile setting for dot1p marking irrespective of the actual packet profile.
- Profile mode queues are supported but with limited stats (only: out-of-profile offered and out-of-profile forwarded).

12.47 Statistics (VSR)

- IP interface statistics are partially supported for SAP and R-VPLS interfaces (egress stats are supported but not ingress).
- On SAPs, the ingress (priority mode) queue statistics do not increment properly for Off.HiPrio (all packets are counted as low priority) and For.InProf (all packets are counted as out of profile).
- On SAPs, the egress queue stats do not increment properly for For.In/InplusProf (all packets are counted as For.Out/ExcProf).
- SDP binding stats do not count ingress dropped packets/octets.
- LSP ingress and egress statistics are not supported.
- The ingress octets traffic statistic for management ports (for example, displayed by the **show port A/1** command) does not increment past zero. [251651]

12.48 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 entries in an existing **match-list** used in a CPM and/or IOM filter may result in a small amount of packet loss in cases where the filter entry references multiple **match-list**s or a **port-range** on FP2- or FP3-based cards. [265287]
- Configuration rollback changing a card **filter-profile** is not supported on 7750 SR-a platforms; each card **filter-profile** in the system must remain the same between two configuration rollback checkpoints. [313843]
- Filter policy of **type src-mac** configured on a VPLS egress endpoint imposes the following restrictions on the QoS policy assigned to the same egress endpoint:
  - **qos ip-criteria** match **dscp** and **src-ip** is not supported and should not be used on a **vpls egress** endpoint if the same egress endpoint is already configured with an **ip-filter** policy of **type src-mac**.
  - **qos ipv6-criteria** match **src-ip** is not supported and should not be used on a **vpls egress** endpoint if the same egress endpoint is already configured with an **ipv6-filter** policy of **type src-mac**.
- Filter policy of **type packet-length** configured on egress endpoints imposes the following restrictions on the QoS policy assigned to the same egress endpoint:
  - **qos ip-criteria** match **dscp** is not supported and should not be used on **egress** endpoints if the same egress endpoint is already configured with an **ip-filter** policy of **type packet-length**.
  - **qos ipv6-criteria** has no restrictions in case of **ipv6-filter** policy of **type packet-length**.

12.49 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 classic 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.

12.50 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 usernames. [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 bound to physical ports require at least IOM3-XP/-B/-C and are supported with the HS-MDAv2
  - Hash and entropy labels are not supported on PW-SAPs bound to physical ports
  - For hardware support and FPE based PW-SAPs see the PXC section of the Usage Notes chapter.
  - 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
  - Mixed SDP types are not supported
- 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:
    - E-Tree
    - 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.
    - 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]
- 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
- 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 causes an error. [265483]
- When **bgp-evpn**>**vxlan** is enabled in Epipe services, the following features are not supported:
  - Endpoints
  - BGP-MH sites
  - Per-service hashing
  - PBB Epipes
  - BGP-VPWS or spoke-SDP FEC
  - **ignore-oper-down**
- The following features are not supported in Release 16.0.R1 along with dual BGP-instance EVPN-VXLAN services:
  - I-VPLS/B-VPLS/M-VPLS/E-Tree service types
  - BGP-VPLS
  - R-VPLS
  - EVPN Ethernet Segment association with **vxlan instance 1** when instance 2 is VXLAN (and not MPLS)
  - SAPs and SDP-bindings when **bgp-instance 1** and **bgp-instance 2** are both associated with VXLAN

  - ESI-based PBR on the VPLS service. ESI-based PBR is supported on single **bgp-evpn vxlan** instance services, and only with **vxlan instance 1**.
  - IGMP/PIM-snooping
- When enabling existing IES features on interfaces linked to EVPN-VXLAN R-VPLS or EVPN-MPLS R-VPLS interfaces, the following commands are not supported:
  - **if**>**vpls**>**evpn-tunnel**
  - **bgp-evpn**>**ip-route-advertisement**
  - **arp-populate**
  - **authentication-policy**

  Dynamic routing protocols, such as IS-IS, RIP, and OSPF, are also not supported.
- The **force-qinq-vc-forwarding** [**c-tag-c-tag | s-tag-c-tag**] command is not supported under **vpls**>**bgp-evpn**>**mpls** or **vpls**>**bgp-evpn**>**vxlan** contexts, and it is not allowed along with the following features:
  - Multi-segment PW
  - Routed, E-Tree, or PBB VPLS services (including B-VPLS and I-VPLS)
  - L2PT termination on the SDP-binding under which the feature is configured
  - IGMP/MLD/PIM snooping within the VPLS service
- 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
  - 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
  - Reassembly for IPv6 as GRE delivery (transport) protocol is not supported

  - Fragmentation for IPv6 as GRE delivery (transport) protocol is not supported
  - L2oGRE with IPv6 transport in a system that contains any FP2-based line cards is not supported
  - L2oGRE packets with IPv6 as the delivery protocol must have the next-header value in the IPv6 header set to “GRE” (47). IPv6 extension headers in the L2oGRE transport IPv6 header are not supported, and packets containing IPv6 extension headers are dropped. **[NEW]**
- Only **bgp-evpn**>**mpls** can be added to an endpoint and only in an Epipe service (no **bgp-evpn**>**vxlan** and not on VPLS services).
- Only a SAP (not part of any endpoint) and a manual spoke-SDP (on the same endpoint) can be configured at most along with **bgp-evpn**>**mpls**.
  - If **bgp-evpn**>**mpls** and the SAP are already configured, the new spoke-SDP must be configured with an endpoint that matches the endpoint configured on **bgp-evpn**>**mpls**.
  - Only one explicit endpoint is allowed per Epipe service with **bgp-evpn** configured.
- A limited endpoint configuration is allowed in Epipes with **bgp-evpn**, in particular, **no active-hold-delay** or **revert-time** are configurable.
- When **bgp-evpn**>**mpls** is added to an explicit endpoint along with a spoke-SDP, the **spoke-sdp**>**precedence** command is not allowed. The spoke-SDP will always have a precedence of 4, and **bgp-evpn**>**mpls** a better precedence.
- VSR does not alias traffic in an EVPN-VPWS all-active multi-homing configuration.
- The association of a PW-port SDP to an ES or vES is not supported. In order to associate a PW-port to an ES, the keyword *pw-port* itself must be used in the **config**>**service**>**system**>**bgp-evpn**>**ethernet-segment**>**pw-port** context. [301998]
- The **bgp-evpn**>**mpls**>**force-vlan-vc-forwarding** command is not supported in BGP-EVPN services when **sap**>**ingress**>**vlan-translation** *vlan-id* is configured. [314136]
- MACsec CLI supports sub-ports. Sub-ports are configured for per-VLAN MACsec configuration. Each sub-port represents a dot1q- or QinQ-tagged traffic. **encap-match** under a sub-port can be used to match a specific dot1q or QinQ traffic flow. Per-VLAN configuration is not supported and only sub-port 1 should be configured. In addition, only **encap-match all-encap** is currently supported.

12.51 EVPN Multicast

- In EVPN OISM Optimized Inter-Subnet Multicast (OISM) scenarios, IGMP-snooping must be enabled on the ordinary R-VPLS services. Without it, multicast traffic is never forwarded to SBD EVPN destinations. [334294]
- In an EVPN OISM where the ingress PE1 sends multicast traffic to an egress PE2 on R-VPLS1, and the receiver is attached to R-VPLS2, if the user performs a **bgp-evpn**>**mpls shutdown** on R-VPLS1, the multicast traffic will arrive at PE2 on the SBD as expected, however it will be dropped. This is because the RPF-check will not pass since the IIF is still pointing at the R-VPLS1.
- The following limitations must be considered when deploying an EVPN OISM solution:
  - The SBD R-VPLS must be configured on single EVPN-MPLS instance R-VPLS, using the default instance and it is the only R-VPLS service for the tenant where **sel-mcast-advertisement** should be configured.
  - The SBD R-VPLS does not support VSI-import policies.
  - This solution is not supported in Routed I-VPLS services.
  - When an R-VPLS is configured as SBD, no unicast routing protocols should be configured under the interface for the SBD (for example, OSPF, IS-IS, VRRP, etc.)
  - All other (ordinary) R-VPLS in the OISM domain (if EVPN-enabled) need to be single instance and using the default instance and EVPN-MPLS only.
  - ECMP is not supported on OISM R-VPLS services for multicast traffic.

12.52 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 SDP to an ES or vES is not supported. In order to associate a PW-port to an ES, the keyword **pw-port** itself must be used in the **config**>**service**>**system**>**bgp-evpn**>**eth-seg pw-port** context.

- 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]
- The following features are not supported when EVPN-VXLAN multi-homing is configured in a VPLS or R-VPLS service:
  - Proxy-ARP/ND
  - EVPN-MPLS instance
  - Association of the VXLAN instance 1 (and **service-id**) to a **network-interconnect-vxlan** ES
  - VXLAN instance 2
  - **source-vtep-security**
  - **vxlan-src-vtep**
  - ESI-based redirection
  - MLD or PIMv6 snooping
  - Multicast data-plane-driven state synchronization based on ESI labels. This is due to the fact that no ESI labels exist in VXLAN encapsulation.
  - Assisted Replication

12.53 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

12.54 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)

12.55 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

12.56 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 DHCP ACK messages, if the SR OS still has an outstanding RADIUS transaction from the first DHCP ACK when receiving the second DHCP ACK, the latter one will be dropped and a DHCP RELEASE message will be incorrectly generated towards the DHCP server. When RADIUS authentication is successful for the first DHCP ACK, the client will still receive the DHCP ACK 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]
- Leaking of a subscriber prefix from a local 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, 139552, 140643]
- 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]
  - 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]
  - 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
  - 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]
- Exporting routes into OSPF that have a subscriber interface as next-hop is not supported. Forwarding for these routes will fail and the following event may be logged when the route is installed: "IOM:find_ip_nexthop Cannot have IP_NEXTHOP on subscriber interface". [196933]
- 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 SLAAC IPv6 hosts that share a /64 prefix (**ipoe-bridged-mode**) with separate **sla-profile** instances are not supported. In 64-bit WAN mode, IPoE DHCPv6 hosts that share a /64 prefix (**ipoe-bridged-mode**) must share the same SLA-profile. Egress traffic for these hosts will share a single (arbitrary) set of SLA-profile instance queues/policers. If 128-bit WAN mode is enabled, IPoE DHCPv6 hosts in **ipoe bridged-mode** can each have different SLA-profiles. [199934]

- The oversubscribed multi-chassis redundancy model in ESM has the following limitations:
  - 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 **anti-spoof 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 database 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.
  - Having both **sla-profile** *sla-profile-name* **host-limits overall** *1* and **host-limits remove-oldest** enabled is not supported. A new host or IP address for the same host cannot bounce the old host in an IPoE session without deleting the IPoE session first.
- 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
- 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 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 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, 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.
- 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)
- In ESM L2-Aware NAT service chaining, when configuring an IPv6 **vxlan-vtep-range**, only SF-IP-only configuration in the forward action of a Value-Added Services (VAS) filter is supported. An SF-IP + ESI configuration in the forward action of a VAS filter is not supported when the **vxlan-vtep-range** is an IPv6 range.
- In ESM L2-Aware NAT service chaining, the VXLAN VNI is ignored when receiving a packet. If the flow exists in the routing context, the packet is forwarded irrespective of whether it was expected to be coming from the Value-Added Services (VAS).
- Data-triggered host setup is not supported in a Wholesale/Retail configuration when the retailer has **private-retail-subnets** configured. [269987]

- An ARP packet is incorrectly not recognized as a valid trigger packet for data-triggered Dynamic Data Services (DDS) authentication when both the DDS capture SAP and a subscriber management capture SAP are configured within the same VPLS service. A workaround is to use a different VPLS service for the DDS capture SAP. [286965]
- SLA-Profile Instance (SPI)-sharing per session or per group of sessions is not supported in the following instances:
  - when **ipoe-session** is disabled on the group interface
  - for static hosts
  - for IPoE or PPPoE sessions active on HS-MDAv2 SAPs
- The SR OS node does not respond to traffic destined to a Wholesale/Retail managed route pointing to a subscriber host with ICMPv4 "Fragmentation Needed and Don't Fragment was Set" or ICMPv6 "Packet Too Big" replies when the DF (Do Not Fragment) flag is set and the negotiated MTU is exceeded.
- The following limitations apply to S-VLAN statistics collection:
  - Not supported on HS-MDAv2
  - Not synchronized across chassis in dual-homing environment
  - Subscriber policer or queue statistics are not synchronized between the forwarding complexes participating in a LAG
  - Policers in stat-mode ‘no-stats’ are not supported. This means that such policers will not collect statistics (count will be 0) even though traffic is flowing through them.
  - If an SLA-profile for the last single host on the S-VLAN is changed (the last SLA-profile instance on the S-VLAN), then the S-VLAN stats will be reset to 0.
  - Dynamic policers (policers created dynamically via Gx/RADIUS) are not supported. If they are used while S-VLAN statistics collection is enabled, their traffic will be double-counted (once through the policers, and once through the next QoS entity that the policer is tied to).
- SAP QoS is not supported for PPPoE traffic redirected to or from an Epipe when **profiled-traffic-only** is configured for the SAP. [293595]
- When GTP-U traffic is received for an unknown TEID on an S11 endpoint, the system does not always send an Error Indication message in response. A counter for these erroneous packets is present in the output of **show subscriber-mgmt gtp statistics**.
- Upon a High-Availability CPM switchover for a redundant BNG pair with Multi-Chassis Synchronization (MCS), SRRP, and a scaled number of subscriber hosts, a few ICMP packets that are destined to the subscriber-interface subnet IP address can be lost. [336377]

- The following limitations apply to Multi-Chassis Support for Usage-Monitoring (Gx):
  - No MISSU support
  - SRRP switchover in combination with an SLA profile change can lead to usage monitoring becoming disabled
  - In scaled scenarios, some instabilities may occur
  - SRRP switchover in combination with Session Level Usage Monitoring is only supported with the default session level monitoring key name "_TMNX_SessionLevelMonitoring"
- For 7750 SR-1s/2s/7s deployments with subscriber management enabled on Link Aggregation Groups (LAGs), the following restrictions apply: [337900]
  - All member ports on an XCM-s must be part of the same Forwarding Plane (FP). Member ports can be on multiple XCM-s cards as long as per XCM-s all member ports are part of the same FP. Use the **show datapath** *slot-number* **detail** command to see the port to FP mapping.
  - Adding ports to or removing ports from a LAG with member ports on the same FP per XCM-s or on multiple XCM-s cards, can take longer than expected when a large number of subscribers are active on the LAG. During that time, the system may not be responsive to subscriber control traffic, such as DHCP renewals. Nokia recommends performing such provisioning actions during a maintenance window. Note that this restriction does not apply for operational state changes of member ports, such as optical link failures.
- The **subscriber-mgmt local-user-db** command gracefully handles configuration collisions by moving the conflicting host entries onto the unmatched host list with a dedicated reason (for example, duplication). Conflicting host entries can be the result of a misconfiguration or a configuration change as part of a transition workflow. The order in which host entries are inserted or moved to the unmatched host list is not preserved after a CPM reconciliation followed by a CPM High-Availability switchover or after a configuration save followed by a node reboot. To prevent unpredictable host lookups, the unmatched host list must be kept empty regardless of the unmatched reason. **[NEW]**
- SAP ingress QoS policies with an **ip-criteria** of **type vxlan-vni** should not be used on subscriber MSAPs. Associating a **sap-ingress** QoS policy with the **group-interface sap** or **msap-policy ies-vprn-only-sap-parameters** command is not supported. **[NEW]**

12.57 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.

12.58 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]

12.59 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 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 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.

12.60 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]

- When routed traffic is forwarded over the IES/VPRN R-VPLS interface that has **vpls**>**discard-unknown** configured, and there is no FDB entry corresponding to the next-hop IP's ARP, frames will be flooded in the VPLS (as opposed to being discarded. [291910]
- Non-system IPv4/IPv6 termination is supported but only in the Base router (not on a VPRN service).
- When IP (IPv4 or IPv6) multicast traffic is forwarded within an R-VPLS service to the associated IP interface but without a Layer-3 OIF instantiated, PIM statistics (“Curr Fwding Rate”, “Forwarded Packets”, and “Forwarded Octets”) in the **show router** *router-instance* **pim group detail** output are incremented. However, MFIB statistics (“Matched Pkts” and “Matched Octets”) in the **show service id** *service-id* **mfib statistics** output are not incremented. [305233]
- A Routed I-VPLS service does not support the use of **provider-tunnel** in the attached B-VPLS service. The CLI commands to set up a routed I-VPLS along with a B-VPLS with **provider-tunnel** are not restricted.

12.61 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]

12.62 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]

12.63 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 CPM 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]

12.64 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, the policy will use the configured priority to affect VRRP instances. Upon a CPM 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 correctly 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]

12.65 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.
- The configuration of a local interface's IP-address as **egr-vtep** *ip-address|ipv6-address* under the **config**>**service**>**epipe|vpls**>**vxlan** context is a misconfiguration, but the system does not block it. [241289]
- 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
- The configuration of the same **egr-vtep** *ip-address*|*ipv6-address* in services using different router instances (Base or VPRN) to terminate VXLANs is not supported, but the system does not block this configuration. [260893]
- Configuring a **vxlan**>**tunnel-termination** *ip-address* with the same *ip-address* as an IP protocol address (for example, Base router VRRP/SRRP backup IP address, **bfd-enable** destination IP address, or redundant interface remote IP) is a misconfiguration, but the system does not block it. [301701]
- Static VXLAN (or a VXLAN-instance with configured **egr-vtep**s) is not supported in conjunction with the following features:
    - I-VPLS/B-VPLS/M-VPLS/E-Tree
    - **allow-ip-int-bind**>**forward-ipv4 | ipv6-multicast-to-ip-int** and **igmp | mld-snooping**
    - ETH-CFM extraction
    - IGMP/MLD/PIM-snooping
    - BGP-AD/BGP-VPLS/BGP-MH
    - Proxy-ARP/ND
    - Endpoints
    - Subscriber-management SAPs
    - MVR
    - Video/management interfaces
    - SDP-bindings in default domain
    - The following FDB related features:
        - **config**>**service**>**vpls**>**mac-protect**
        - **mac-move**
        - **mac-pinning** on SAP/SDP-binds along with static VXLAN

  - **restrict-unprotected-dst**
  - Assisted-Replication in the static VXLAN-instance (AR requires EVPN control plane)
  - STP, BPDU-translation, or L2PT
- A static **vxlan-instance 1** can be associated to a **network-interconnect-vxlan** Ethernet Segment (ES), as long as **bgp-evpn mpls** is also enabled in the same VPLS service. The behavior of the NDF with respect to the BUM traffic received on **vxlan-instance 1** is governed by the new command **rx-discard-on-ndf** {**bm** | **bum** | **none**}.

12.66 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 ETH-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.

12.67 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]

12.68 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]

12.69 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.
- 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]

12.70 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.
- The **oam vccv-ping** command 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 (**configure li li-source** *mirror-svc* **nat**) 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]
- Egress MPLS traffic on an interface is not mirrored. This limitation applies to VSR.
- Mirroring/LI is not supported for LNS subscribers. This limitation applies to VSR.
- LI log events (for example, as syslog or SNMP notifications) cannot be sent out of a VPRN interface. The configuration of "from li" (or "source li" in MD interfaces) is not supported in **configure service vprn log log-id**. The NSP cannot manage LI on SR OS routers via a VPRN.

12.71 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.

12.72 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**>**servers**>**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).
- Dynamic ports are always reserved, even if only deterministic port blocks have been reserved via configuration. [195357]
- 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]
- 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.

12.73 NAT (VSR)

- **sap-ingress fc in-remark** and **sap-ingress fc out-remark** are not being applied for L2-Aware NAT subscriber traffic. [268570]

12.74 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.

12.75 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 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]
- 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]
- 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

- AA captive HTTP-redirect cannot be used in ESM, WLAN-GW ESM and vRGW subscriber deployments where L2-Aware NAT is configured. Large-scale NAT may be used as an alternative in all three deployment models with AA captive HTTP-redirect. Alternatively for WLAN-GW, DSM subscribers may be used with AA captive HTTP-redirect and L2-Aware NAT. [260531]
- In HTTP Header Certificate-based Encryption, no enrichment takes place if the resulting packet size will exceed the configured MTU size. The length of an encrypted header is directly proportional to the length of the encryption key: the longer the encryption key, the longer the encrypted header. Operators should therefore be cautious when defining the key length and selecting which fields will be encrypted and enriched.
- The following features cannot be used when AA Layer-7 classification is disabled. [287041]
  - HTTP/HTTPS Header enrichment
  - Application start/stop notification
  - HTTP Redirect, HTTPS Redirect, Error-Redirect, Notification
  - **dns-ip-cache** for rating group anti-fraud verification
  - ALGs (flow wildcarding)
  - GTP/SCTP Firewall
  - **url-filter** (parental control via ICAP or local list)
  - tethering detection
- AARP remote node shunts are not supported on imm12-10gb-sf+, imm1-100gb-cfp, and imm3-40gb-qsfp cards. [304260]
- Synchronization of Usage-Monitoring counters in Gx between chassis is not supported for AA.

12.76 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]

- 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:
  - IS-IS protocol traffic
  - BFD over LAG link members (uBFD)
  - Layer-2 control traffic in IP interface
- Cflowd sampling of egress multicast traffic on a subscriber management group interface will only sample traffic that uses a multicast MAC address to send the traffic. If a unicast MAC address is used to send the multicast traffic then that traffic will be sampled as unicast. [304585]

12.77 sFlow

- In Releases 13.0.R6 and higher, 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]
- sFlow is not supported on satellite ports.

12.78 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 does not support routes tunneled over LDP shortcuts (**configure router ldp-shortcut**). [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 High-Availability switchover; it is expected they restart from zero after the switchover. [250631]
- For Seamless-BFD (S-BFD) over SR-TE LSPs, the system will always set the destination address of the S-BFD packet at the initiator to the IP address of the reflector. The reflector will set the destination address of the reflected packet to the source address of the received S-BFD packet and will set the source address of the reflected packet to the destination address of the received S-BFD packet. An S-BFD packet received with a 127/8 destination address will be dropped.

12.79 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, or 7950 XRS as the Customer Edge (CE) hub. As a workaround, the 7750 SR or 7450 ESS, 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 svc-ping** command on a VPRN service with the **local-sdp** and **remote-sdp** keywords configured is always sent over the control plane. It is better to use the **oam vprn-ping** command instead. [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.
- When SAA ETH-CFM continuous tests are configured and CPM-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]
- The **oam vprn-ping** command 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 **oam vxlan-ping** command will accept untagged or single VLAN-tagged packets with Ethernet-type 0x8100. Any VLAN-tagged packet that uses anything other than Ethernet-type 0x8100, or any packet with more than a single VLAN tag, will not be processed by the **vxlan-ping** function. [294786]
- The Hashing Visibility Tool (**find-egress**) does not support test packets to multicast or broadcast addresses. Test packets with multicast or broadcast addresses will be discarded without determining an egress port.
- The **find-egress** command cannot be used with the following port types as the ingress port:
  - ATM
  - POS
  - Satellite ingress
  - Forwarding that includes a PXC

12.80 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.

12.81 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.

12.82 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 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.
- OF-switch does not support selection on out_group in flowmod message for delete and delete_strict. [326844]

12.83 NETCONF

- The NETCONF interface (with Nokia SR OS YANG models) is a model-driven (MD) interface and is subject to the general limitations of MD interfaces. See the Model-driven Interfaces section of Known Limitations for more information.
- Do not deploy NETCONF in a production network without carefully reviewing the Known Limitations and consulting with your Nokia representative.
- The classic 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 classic 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 Nokia SR OS YANG modules in the NETCONF interface:
  - SR OS does not support logging of operator actions (that is, who issued which commands) via NETCONF with the Nokia SR OS YANG modules.
- The following items apply to the use of the Base-R13 SR OS YANG modules and data model in the NETCONF interface (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 classic CLI infrastructure. That linkage results in NETCONF behavior that follows many classic 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 the classic 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>).
- Leafs of type leaf-list are matched only if "content match nodes" for all its values in the correct order are present within a single containment node.
- The NETCONF login/logout events (and event counters) differ depending on which port is used to access the SR OS NETCONF server. [334021]

  Using port 830:

    - When a NETCONF user login/logout successfully, the following events are generated:
        - USER/**netconf**_user_login
        - SECURITY/**netconf**_user_login
        - USER/**netconf**_user_logout
        - SECURITY/**netconf**_user_logout

- When a NETCONF user login unsuccessfully, the following events are generated:
  - USER/**netconf**_user_login_failed
  - SECURITY/**netconf**_user_login_failed
  - USER/**netconf**_user_login_max_attempts
  - SECURITY/**netconf**_user_login_max_attempts

Using port 22:

- When a NETCONF user login/logout is successful, the following events are generated:
  - USER/**cli**_user_login
  - SECURITY/**ssh**_user_login
  - USER/**cli**_user_logout
  - SECURITY/**ssh**_user_logout
- When a NETCONF user login in unsuccessful, the following events are generated:
  - USER/**cli**_user_login_failed
  - SECURITY/**ssh**_user_login_failed
  - USER/**cli**_user_login_max_attempts
  - SECURITY/**ssh**_user_login_max_attempts

12.84 ISSU

- ISSU can use the Soft Reset mechanism and if used, is subject to any limitations of Soft Reset in the source/starting release of the upgrade. See Soft Reset in the Known Limitations section for the source/starting release.
- Limitations specific to ISSU across minor releases (“Minor ISSU”) are as follows:
  - Minor ISSU refers to ISSU across maintenance releases (e.g. from 19.10.R1 to 19.10.R2) and not to ISSU across minor releases (e.g. from 19.7.R2 to 19.10.R2, which is not supported. For releases 19.10 and later, minor ISSU is only supported across maintenance releases that belong to the last minor release of the year, which for 2019 is 19.10.
  - Minor ISSU is supported on platforms with redundant CPMs.
  - Minor ISSU is supported across up to a maximum of 20 minor releases.
- Limitations specific to ISSU across major releases (“Major ISSU” or “MISSU”) are as follows.

  - For releases 19.10 and later, MISSU is only supported to and from a maintenance release that belongs to the last minor release of the year, which for 2019 is 19.10.
  - MISSU is supported on platforms with redundant CPMs.
  - MISSU is supported across two major releases (i.e., Release 15.0 to Release 19.10) for all paths 15.0.Ra → 19.10.Rb where:
    - a is >= 4 and b is >=1
    - The release date of 19.10.Rb is at least 90 days later than the release date of 15.0.Ra.
  - MISSU is supported across a single major releases (i.e. Release 16.0 to Release 19.10) for all paths 16.0.Rc -> 19.10.Rd where:
    - c is >= 4 and d is >= 1
    - The release date of 19.10.Rd is at least 90 days later than the release date of 16.0.Rc.
  - A MISSU switchover, when a multi-chassis APS port is active and the VRRP port feeding that APS port is master as well, may result in a longer outage on impacted channels.

    As a workaround, either the APS ports or the VRRP master should be moved to the other MC-APS router before the MISSU upgrade. [157196]
- In MC-IPsec scenarios, a multi-chassis switchover to the standby chassis must be performed before performing ISSU; otherwise, an extended data loss may occur if the MCS link goes operationally down.
- A mandatory firmware upgrade on an MDA/XMA/IMM will cause a hard reset (instead of being able to Soft Reset). A Deferred MDA Reset is not supported for these cases. A hard reset must be performed during ISSU if the starting release is earlier than a mandatory firmware upgrade and the target release is equal to or later than the firmware upgrade.

  Mandatory firmware upgrades apply to the following cards and releases:
  - 15.0.R4: x40-10g-sfp XMA and imm40-10gb-sfp IMM [255711]
  - 16.0.R4: me1-100gb-cfp2 MDA [279794]
  - 16.0.R4: p10-10g-sfp, p6-10g-sfp, and p1-100g-cfp MDAs [299643]
  - 16.0.R4: cx20-10g-sfp and cx2-100g-cfp XMAs [299643]
  - 16.0.R6: me10-10gb-sfp+ and me6-10gb-sfp+ MDAs [307999]
  - 16.0.R7: imm40-10g-sfp IMM and x40-10g-sfp XMA [303843]
  - 19.10.R1: me2-100gb-cfp4 and me2-100gb-qsfp28 MDAs [251111]
- A non-mandatory firmware upgrade on an MDA/XMA/IMM will not cause a hard reset. The card will Soft Reset instead. A deferred MDA reset is supported for these cases. The firmware upgrades apply to the following cards and releases:
  - 19.10.R3: me10-10gb-sfp+ and me6-10gb-sfp+ MDAs [338476]

12.85 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]

12.86 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.

12.87 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]

12.88 WLAN-GW

- The distributed RADIUS proxy is only guaranteed to handle Access-Request packets of up to 1024 bytes. [221041, 241114]
- WLAN-GW Layer-2 AP SAPs do not support **connection-profile-vlan**. In WLAN-GW Layer-2 wholesale context, SAPs with **connection-profile-vlan** are not supported in the retail VPLS where the WLAN-GW node is configured. [304528, 304794]
- VLAN-aware distributed RADIUS proxy is not supported in combination with Layer-2 AP tunnels. [335236]

12.89 VSR Infrastructure

- When SR-IOV is used a change to the host interface MTU disrupts traffic flow to and from the virtual machine.
- VM snapshots are not supported with VSR virtual machines.
- Live migration is not supported with VSR virtual machines.
- A VSR system may become unstable or reset if one of its vCPUs is pinned to a host CPU core that is also servicing a high rate of IRQs associated with inbound VirtIO traffic (note that the isolcpus kernel parameter does not affect the scheduling of IRQs).
- The OVF file provided for VSR deployment on VMware ESXi hosts does not include a vApp section. The parameters described in the vApp deployment section of the *VSR Installation and Setup Guide* must be set manually, after deployment, using vCenter (Advanced Settings).
- When SR-IOV is used with supported Intel NICs only untagged frames are accepted from or delivered to the guest. VLAN tagged frames with ethertype 0x8100 cannot be passed.
- When SR-IOV is used with supported Intel NICs frames transmitted by the guest may be dropped in the NIC if they do not have the source MAC address that is specified in the configuration of the vNIC interface. This anti-spoofing can be disabled using the **ip link set dev** *dev-name* **vf** *vf-number* **spoofchk off** command.
- When SR-IOV is used with Intel X710 NICs the “trust” setting should be enabled for each VF using the **ip link set** *dev-name* **vf** *vf-number* **trust on** command. Note that this is only supported in Red Hat Linux 7.3 or higher. The **trust on** setting is necessary to allow untagged multicast frames to be accepted from and delivered to the guest.
- When SR-IOV is used with supported Mellanox Connect-X4 and ConnectX-5 NICs you must enable the “trust” setting for each VF using a command such as the following: **echo ON** > /**sys**/**bus**/**pci**/**devices**/*domain:bus:device.function*/ **sriov**/*vf_num*/**trust**.
- The packet forwarding rate for traffic transmitted out a Mellanox ConnectX4 NIC port drops substantially when the traffic has a mix of different 802.1p bit settings. This applies to both SR-IOV and PCI pass-through. [This is tracked by Mellanox case number 00311370.]
- The maximum PPS throughput of the Mellanox Connect X4 NIC drops with larger packet sizes. This applies to SR-IOV and PCI pass-through.
- The Mellanox MCX4121A NIC drops QinQ frames transmitted by the guest.

- When SR-IOV or PCI pass-through is used, access to the NIC option ROM should be disabled by a <rom bar='off'/> element in the libvirt XML configuration of the interface. If the option ROM is not blocked, it can result in delays and degraded performance in the VSR guest.
- The Mellanox MCX 416A-CCAT NIC may stop forwarding Layer 2 and multicast traffic after an extended period of inactivity.
- The default speed of all SR OS I/O ports (of the m20-v MDA) is 40 Gb/s; however, the speed can be changed using the **config**>**port**>**ethernet**>**speed** command. Allowable values are 1, 10, 25, 40, 50, and 100 Gb/s. Autonegotiation of the speed with the remote end is not supported. Note that the operational speed affects only the rate of egress traffic on the port, not the rate of ingress traffic.
- The default and maximum MTU of SR OS I/O ports (associated with the m20-v MDA) of a VSR-I system is 9212 bytes.
- Software upgrades using the **admin save detail** command are not supported on the VSR platform. [320203]

12.90 MAP-T (VSR)

- MAP-T is not supported in combination with R-VPLS or L2TPv2/v3. However, this is not blocked in CLI.
- End-User-Prefix (IA-PD) length must be 64 bits or less.
- Rule IPv6 prefix must be of the same length in a MAP-T domain.
- Overlapping IPv6 rule prefixes within a routing context will be blocked in CLI.
- Overlapping IPv4 rule prefixes within a routing context will be blocked in CLI.
- To comply with section 4 of RFC 1191, **ping** to CE MAP address from Border Router (BR) is not supported. However, IPv6 WAN address of the CE can be **ping**ed.
- Rules must be **shutdown** before they are modified.
- IPv6 packets with hop-by-hop (HOPOPT) next header are discarded since there is no equivalent option in IPv4.
- MAP-T function cannot be collocated in the same node with BNG function serving the same customers.
- A MAP domain template can be instantiated inside a single routing context. In other words, the same MAP domain template cannot be applied to multiple inside routing contexts.
- ICMPv4 Error Code 4

- In case that IPv6 MTU in the MAP-T domain (**config**>**service**>**nat**>**map-domain**>**mtu**) is violated by a translated-IPv4 IPv6 packet with DF=1, an ICMPv4 code 4 with 'MTU of next hop' is returned to the sender.
- To comply with section 4 of RFC 1191, VSR Border Router (BR) sends "MTU of next hop" adjusted to a value that is 20 bytes smaller than the configured IPv6 MTU in the MAP-T domain.

12.91 PW-ports Bound to Physical Ports

- 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
  - Hash labels
  - Entropy labels

12.92 FPE-based PW-ports

- 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
  - MC-LAG
  - **force-qinq-vc-forwarding** and **force-vlan-vc-forwarding** functionality. (Note that **vlan-vc-tag** *vlan-id* functionality can be used instead of **force-vlan-vc-forwarding** when the spoke-SDP attached to the PW-port is of type VLAN)
  - Traffic is forwarded over a L2oGRE tunnel that is terminated on the local subnet (subnet that is directly connected to the VSR node in Base router), only if there is a static route for the far-end address configured

12.93 Flex PW-ports (VSR)

- Traffic is forwarded over a L2oGRE tunnel that is terminated on the local subnet (subnet that is directly connected to the VSR node in Base router), only if there is a static route for the far-end address configured.
- VXLANv4 termination is supported only on system IP address.
- Model-driven management infrastructure (the MD-CLI, NETCONF, gRPC) is not supported for Flex PW-port.
- Queue-groups on PW-port are not supported.
- Only PW-SAP of type **capture-sap** is supported in VPLS.
- The port to which an SDP associated with the PW-port is bound must be configured as a hybrid port. A PW-port associated with an SDP bound to a port which is configured as network will not forward traffic even though the PW-port operational state may be up.
- Flex PW-ports are not supported in a system that has network interfaces configured in a LAG. Consequently, the **config**>**service**>**epipe**>**load-balancing**>**per-service-hashing** CLI configuration has no effect in an Epipe associated with the Flex PW-port.
- ECMP for Flex PW-ports is not supported. Spraying of traffic within PW-port over multiple equal cost paths is not supported. However, multiple PW-ports can be distributed over equal cost multiple paths where each Flex PW-port selects one of the paths. When this path fails, Flex PW-port will be rerouted to one of the remaining paths.
- Endpoint configuration is not supported in an Epipe associated with the Flex PW-port. The **config**>**service**>**epipe**>**endpoint** *endpoint-name* CLI configuration will have no effect.
- **force-qinq-vc-forwarding** in spoke-SDP with Flex PW-port is not supported.

12.94 NGE (VSR)

- When NGE is enabled on a VPRN with auto-bind and packets are fragmented by the VSR before being encrypted and sent across the VPRN network, packets are dropped at the far-end with the reason "Authentication Failure". [274888]

13 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 16.0.R2 to 16.0.R7 also apply to Release 19.*x*. Refer to the most recent *SR OS 16.0 Release Notes* for the summary of resolved issues in Releases 16.0.R2 through 16.0.R7.

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.

13.1 Release 19.10.R6

13.1.1 Hardware

- A port on an FP4 XMA no longer disables the transmitter when recovering from a disturbance (for example, LOS) detected at the receiver. [317543-MA]
- If the port speed is changed from 1G to 10G on MACsec-enabled ports on an me12-10/1gb-sfp+ MDA, frames do not get encrypted. A workaround is a **shutdown**/**no shutdown** of the MACsec sub-port. [346743-MI]
- Provisioning of a license level upgrade against the XMA of a 7750 SR-1s where the upgrade requires a reset of the XMA incorrectly triggers an immediate reset of the node. As a workaround, license level upgrades can be added to the configuration file manually, followed by a manual node reboot. [353044-MA]

13.1.2 Classic CLI

- When a new RSVP point-to-point LSP is automatically established by the auto-LSP feature, the configuration change indicator “*” appears in the CLI prompt even though no configuration changes were made. [344048-MI]

13.1.3 System

- For a RADIUS server connected to a VPRN that is making use of **grt-lookup allow-local-management** to authenticate or account Base router users, the RADIUS reply Access-Accept or Accounting-Reply messages are incorrectly dropped. [341451-MI]
- A default route of 0.0.0.0/0 configured in the BOF is not installed in the management routing instance after a system reboot. [343730-MI]
- MACsec traffic may be disrupted for a number of seconds if a card Soft Reset (**clear**>**card**>**soft**) is attempted while MACsec is configured. Also, MACsec statistics may reset during a card Soft Reset. [344964-MA]
- With PTP enabled, ISSU synchronization from a prior release to Release 19.10.R1 or later may trigger a reset of the standby CPM running a newer software release, which prevents ISSU from completing. As a workaround, disable PTP before starting the ISSU and then re-enable it after the switchover to the new software release. [347684-MI]
- When SSH with **preserve-key** is enabled, after a reboot, host keys may not be recognized and will not be used anymore. The user is prompted to store a host key when connecting to a known server. When the user is prompted to accept an SSH connection, the user should ensure to check the SSH finger print and make sure there is no man-in-the-middle attack being performed. [350273-MI]

13.1.4 Model-driven Interfaces

13.1.4.1 Common Issues

This section lists resolved issues that apply to all MD interfaces.

13.1.4.1.1 System

- State information for **state chassis** *chassis-class chassis-number* is only returned when **configure chassis** *chassis-class chassis-number* is configured. [350167-MI]

13.1.4.1.2 QoS

- If a **scheduler-policy** with scheduler **overrides** is configured under a **configure port ethernet access ingress queue-group** and the **scheduler-policy** is removed and replaced in the same transaction with a different **scheduler-policy** that is compatible with the configured **overrides**, the **commit** fails even though the result is a valid configuration. [318810-MI]

13.1.4.2 MD-CLI

This section lists resolved issues that apply only to the MD-CLI.

- Command completion for references does not operate correctly in the MD-CLI global operational commands, for example **ping router-instance** <TAB>. [318462-MI]

13.1.5 Telemetry

- Mixing the scale paths and non-scale paths within a single subscription message is not allowed; it can result in the “Max number of paths for all subscriptions reached” error message. [352413-MI]

13.1.6 APS

- If a classic CLI **rollback revert** operation must remove or alter the working bundle associated with a BPGrp, then it also deletes and rebuilds any APS port associated with that BPGrp. [121024-MI]

13.1.7 IS-IS

- In specific multi-instance IS-IS scenarios using **loopfree-alternates** and inter-instance export policies, a failure of multiple links may trigger incorrect IS-IS updates, which may lead to incorrect route propagation and possible routing loops. [350994-MA]

13.1.8 BGP

- If the RPKI **connection-retry** is configured with the default of 120 seconds, and there is a connection failure for 2 intervals (240 seconds), the timer does not reset until after the **stale-time** expires (3600 seconds). A workaround is to configure the **connection-retry** to 100 seconds. [334059-MI]
- Changing export policies to include conditional-expression statements causes an increase in CPM memory utilization. [349132-MA]
- RPKI route origin validation updates can cause periods of high BGP CPU usage. During this time there may be issues related to slow BGP CLI display command output, slow BGP route update propagation, and the bringing up of new BGP neighbors. These issues may be more noticeable when there are a scaled number of routes in the RIB-IN and there are frequent changes to the RPKI database. [350708-MA]

13.1.9 MPLS/RSVP

- Using **isis advertise-tunnel-link** in conjunction with **mpls interface srlg-group** results in an invalid IS-IS TE-SRLG TLV for advertised tunnels. This can result in failing or incorrect CSPF calculations. [349048-MA]

13.1.10 QoS

- The SNMP MIB variables in the TIMETRA-PORT-MIB and TIMETRA-SAP-MIB for the queue depth monitoring average elapsed time are incorrectly using seconds instead of ticks. [352147-MI]

13.1.11 Services General

- If the **cpm-http-redirect optimized-mode** command is enabled (default), the active CPM may reset during a web redirect. [347922-MA]
- The **cpm-http-redirect optimized-mode** command does not function correctly in combination with PW ports (FPE-based and non-FPE-based). [350903-MA]

13.1.12 EVPN Multicast

- MCS-synchronized IGMP-snooping entries may be incorrectly deleted in an all-active Ethernet-Segment SAP in EVPN services, when there is a Designated Forwarder (DF) switchover. [345282-MA]

13.1.13 Subscriber Management

- Using a mix of implicit (**msap-defaults** on the **capture-sap**) and explicit (RADIUS or LUDB **msap-defaults**) MSAP parameters is incorrectly allowed for IPoE sessions but it is correctly checked for all other host types. [351007-MI

13.1.14 IMPM

- When a major ISSU upgrade is performed while using IMPM with **round-robin-inactive-records** configured, an additional traffic loss can occur when the standby CPM becomes active because all streams are started on the secondary paths instead of being placed on the available paths in a round-robin order. [347608-MA]

13.1.15 NAT

- Resources counting on Large Scale NAT (LSN) subscribers may be incorrect if the number of subscribers created by Static Port Forwards (SPF) exceeds the maximum. [336681-MI]

13.1.16 Application Assurance

- If the first packet of an RTP or STUN pair arrives in the **to-sub** direction and the next packet arrives in the **from-sub** direction within the flow timeout period and is unable to allocate a flow record resource, the AA ISA card may reset. Flow record resources may not be allocated under overload conditions or when it is blocked by a flow record policer, flow rate policer, or session filter. [350439-MI]

13.1.17 BFD

- The output of the **show**>**router**>**bfd**>**session** command for sessions with IPv6 link-local addresses has been changed. Previously, the interface name was displayed as part of the destination address. The output has been updated to display the source address, with the interface name as part of that source address for the session. [351912-MI]

13.1.18 VSR

- Executing the **ping** command with a certain packet size could result in a VSR reset if the SMBIOS vsr-deployment-model option is set to queue-scale and if the ping packet is transmitted over a certain NIC type. [337858-MA]
- A VSR using the Mellanox Connect-X-5 NIC does not boot up when issuing the **admin**>**reboot** command or during a system reset. [351699-MA]

13.1.19 Issues Resolved in Prior Releases

See the item marked **[NEW]** in Release 19.10.R3.

13.2 Release 19.10.R5

13.2.1 Hardware

- On a 7950 XRS system, ingress “FCS Errors Detected” alarms displayed in the **show**>**card**>**detail** output are incorrectly not erased after the XMA card reporting the errors is reset. [285822-MI]
- After an ISSU or soft reset of an MDA, the Transceiver Digital Diagnostic Monitoring (DDM) threshold crossing event logs might not get generated. A workaround is to perform a port **shutdown** followed by **no shutdown.** [329835-MA]
- On a 7750 SR-2s chassis, the output for **show**>**sfm**>**detail** displays all zero values under Hardware Resources. This is expected as the SFM power is accounted for in the XCM numbers on the SR-2s but the zero power values should be removed from the **show**>**sfm**>**detail** command output. [340466-MI]

13.2.2 Classic CLI

- The log event "CRITICAL: LOGGER #2002 Base A:RAPIDLOGIC:UNUSUAL_ERROR "Slot A: RCC_OM_WriteHandle" is generated when events are being written to a log that is configured with **to cli** and a user that is subscribed to that log enters a CLI command using an output modifier (a **match** command), and the output of the CLI command is waiting for user input. This log event is benign. [345340-MI]

13.2.3 System

- Log filters do not function correctly in the **li**>**log** context. [331417-MI]
- A node reboot may occur if a user configures an invalid management IPv6 address followed by a valid IPv6 address in the BOF. The workaround is to remove the invalid management IPv6 address first, before adding the valid IPv6 address. [341591-MA]
- Management access from a VPRN via GRT leaking to an address in the base instance does not work unless the **config**>**system**>**security**>**management**>**allow-ssh** command is enabled. [342664-MI]

13.2.4 MD-CLI

- Unnamed parameters with quoted values in show commands are not displayed in command completion and complete the next parameter instead. For example, when using the **show log event-parameters** command, the **event-name** parameter is not displayed and completes the **event-name** parameter instead. [338905-MI]
- The **tree flat detail** command displays duplicate types. [341839-MI]
- In mixed-management interface configuration mode, a user profile cannot be configured to only use the MD-CLI. [343843-MI]

13.2.5 NETCONF

- An error occurs when a list's key node is used as both a “content match node” and a “selection node” in the same request. [331380-MA]
- The combined nokia-conf/nokia-state revision dates returned in reply to a get/filter/netconf-state/schemas request, may not match the actual revision dates in the combined nokia-conf/nokia-state YANG files. [344308-MA]

13.2.6 Model-driven Interfaces

13.2.6.1 System

- In model-driven configuration mode, when a system fails to bootup either due to a problem while loading the configuration file or due to a commit failure, the system shuts down the SNMP and the candidate datastore may become out of sync. [335100-MI]

13.2.6.2 QoS

- If the classic CLI command **adv-config-policy**>**child-control**>**offered-measurement**>**add** *percent* is configured with a non-zero value, followed by a non-default value for **adv-config-policy**>**child-control**>**bandwidth-distribution**>**above-offered-allowance**>**unconsumed-higher-tier-rate** <*percent*>, the value of the **add** *percent* in model-driven interfaces is zero. Also, if the classic CLI command **adv-config-policy**>**child-control**>**offered-measurement**>**granularity** *percent* is configured with a non-zero value, followed by a non-default value for **adv-config-policy**>**child-control**>**bandwidth-distribution**>**above-offered-allowance**>**delta-consumed-higher-tier-rate** *percent*, the value of the **granularity percent** in model-driven interfaces is zero. [345778-MI]

13.2.6.3 Routing

- Route-tag "0" is incorrectly allowed in the model-driven CLI when configuring the route-tag value in **service**>**ies**|**vprn**>**interface**>**ipv4**|**ipv6**>**neighbor-discovery**>**host-route**>**populate** {**static**|**dynamic**|**evpn**} **route-tag** <*number*>. This "0" value is not supported and should not be configured. [344313, 344791-MI]

13.2.6.4 Ingress Multicast Path Management

- The **configure multicast-management chassis-level per-mcast-plane-capacity total-capacity** element in MD interfaces accepted a value of 16500 from Release 16.0.R2 until 19.10.R1, although it was not a valid option. In Release 19.10.R1, the value was removed from the YANG models and MD interfaces without deprecation or obsoletion. In Release 19.10.R5, the value is added back into the YANG models and marked obsolete. During an upgrade in MD configuration mode to Release 19.10.R5 or higher, a value of 16500 is accepted but then converted into the value 15000. [343076-MI]

13.2.7 LAG

- Adding a port to a LAG with **pw-port** bindings from a slot or MDA that already has one port in that same LAG might fail if some of the slot or MDA allocated resources used by all **pw-port** SAPs and subscribers are greater than the free resources.

  The following resources are impacted:

  - Slot resources:
    - QoS user schedulers
    - QoS user scheduler overrides
    - Subscriber SLA profile instance

        - Subscriber SLA profile instances QoS overrides
        - HSMDA queue overrides
        - QoS intermediate arbiter overrides
    - MDA resources:
        - Egress HSMDA queue groups (as seen in the output of **tools dump resource-usage card** *slot-num* **mda** *mda-slot*) [343415-MI]
- After a 7750 SR-s reboot, the active MC-LAG state sometimes reverts. A workaround is to configure an access port **hold-time up** timer. [343511-MI]

13.2.8 PTP

- The 7750 SR-7s/14s CPM LED mappings are incorrect resulting in the OES2 link and activity lights being lit when the SyncE/1588 port is used. [340195-MI]

13.2.9 Routing Policies

- If a new community name is added, and an **admin**>**save** is done prior to the policy changes being committed, all previously configured communities are removed from the configuration file. Another **admin**>**save** is required after the policy changes have been committed to have the communities added back to the configuration file. [344970-MI]

13.2.10 OSPF

- After a rollback from a configuration that contains OSPF multi-instance to a configuration that does not contain OSPF multi-instance, the **ospf multi-instance** command is still present. The workaround is to remove the **ospf multi-instance** command manually before rolling back to the configuration that does not contain the **ospf multi-instance** command. [334353-MA]

13.2.11 BGP

- Receiving a serial notify PDU at the same time as the refresh interval expires may result in receiving duplicate information from the cache server. This causes the RPKI-router session to reset. [331146-MI]

- The node may reset if it receives a remote BGP L2-VPN route that is preferred over an overlapping local route. [341958-MA]
- If the classic CLI **admin**>**rollback**>**revert** command configuration removes the **epipe**>**bgp**>**vsi-export** statement and associated policy from the configuration, this can fail with the error message: “INFO: PLCY #1001 Configuration failed because of inconsistent values - Statement “vsi-export”: Unable to delete because statement is referenced!”. [343678-MA]
- A **router aggregate** with **discard-component-communities** enabled does not always correctly suppress the component communities. [344262-MA]
- Both the active and standby CPM may reset when the following conditions are present: [344513-MA]
    - The **use-bgp-routes** command is enabled
    - Unlabeled and labeled address families are enabled
    - A labeled prefix matches an unlabeled next-hop

13.2.12 QoS

- The configuration of the **scheduler-override** command on Ethernet satellite (**esat**) ports does not take effect when a new resilient port is added and results in a user scheduler resource mismatch within the system. This could prevent other overrides from being instantiated. [303546-MA]
- The rate of a policer, configured with a **percent-rate** and a **scheduler-parent** or **port-parent**, is not updated when its parent rate changes. [345849-MI]

13.2.13 Subscriber Management

- V6 host setup (IPoE and PPPoE) is not supported when the **delegated-prefix-length** is set to **variable** and the DHCPv6 **user-ident** is set to **mac-interface-id.** [332743-MA]
- An ISSU upgrade fails on a system where the subscriber hosts have PCC rules with QoS actions that cause QoS policy clones to be created. The base policy from which these QoS clones have been created must contain IPv4 or IPv6 criteria. [337152-MA]
- Static IPv6 hosts should not copy the default **inter-dest-id** (**def-inter-dest-id**) to their **inter-dest-id** setting when activated. [341973-MA]

- Subscriber hosts connected using redundant Ethernet satellite ports with resilient uplinks terminated on different forwarding complexes of the same FP4 line card, can result in a line card reset upon **system**>**satellite**>**eth-sat**>**port-map** configuration changes. As a workaround, the primary and secondary links should be either terminated on the same forwarding complex of the same FP4 line card or use different line cards. [343829-MA]
- When using IPv6 static forwards for a vRGW IPv6 firewall, the router may become unresponsive when a user accesses tmnxNatFwd2Entry or executes the **show**>**service**>**nat**>**port-fwd-entries** command. [344716-MI]
- Enabling the **control-word** on a **pw-port binding** causes PPPoE LCP Echo packets received from PPPoE clients to be incorrectly processed and PPPoE sessions to be disconnected. [345435-MA]
- When a create-session-request message is received for an IMSI and APN pair for which a PDN session already exists, the old PDN session is deleted. If this PDN delete happens during Local Address Assignment (LAA) of the previous create-session-request message as a result of load, a slow RADIUS response, or LAA, a High-Availability CPM switchover or standby CPM reset can occur. [347440-MA]
- For ESM-over-GTP, after a High-Availability switchover the active CPM can have an IPoE session without an associated GTP session. In such cases, the standby CPM fails to reconcile and the “IPOE_SESSION:UNUSUAL_ERROR ipoeSessionRedUpdateGtpSessionInfo: GTPSession not found to link to session” log event is generated. [348233-MI]

13.2.14 IPsec

- An ISA card may reset while re-assembling IPsec tunneled frames which are larger than 9000 Bytes. [341551-MA]

13.2.15 VSR

- When a VPLS needs to flood frames on multiple SDPs that have NGE enabled, frames get encrypted twice on all but the first SDP. This results in malformed packets getting forwarded to the far-end. [345483-MA]

13.3 Release 19.10.R4

13.3.1 BGP

- When RSVP tunnels are used to resolve VPRN routes, changes in the underlying core network's IGP topology may cause the VPRN route age to refresh unnecessarily, even though the IGP topology change does not directly affect the route's next-hop. This behavior has been present since Release 15.0.R1 and does not affect traffic. [343340-MI]
- A **router aggregate** with **discard-component-communities** enabled does not correctly suppress the component communities when the BGP instance is **shutdown** followed by a **no shutdown**. [345308-MA]

13.3.2 MPLS/RSVP

- When interoperating with other vendors, where the other vendor advertises link local/remote identifiers for interfaces that also have an IPv4 address assigned, it is possible an error may occur after the MPLS **resignal-timer** has expired "CRITICAL: LOGGER #2002 Base A:TE:UNUSUAL_ERROR "Slot A:teDbLnkWithLocalIdGet". In most cases this error is benign.

  For RSVP-TE signaled LSPs there is no impact, except when the cost or the **least-fill** criterion of the old path changes. If the cost of the old path increases and there is another, better cost path with a cost that is equal or worse than the original old path cost, the LSP will incorrectly not be re-signaled.

  If the cost of the old path decreases but it remains the best cost path at that moment, the path is re-signaled unnecessarily.

  If **least-fill** is enabled and the cost of the new path remains the same while the **least-fill** criterion of the old path changes, MPLS may not generate a trap to indicate that a better **least-fill** path is available for the LSP.

  For SR-TE LSPs, the new path is re-programmed unnecessarily, even when the old path is still the best path. [344376-MI]

13.3.3 OAM

- The DSCP of a locally-generated LSP self-ping packet is not set to the expected value of NC1/CS6. [347028-MA]

13.4 Release 19.10.R3

13.4.1 Hardware

- When changing the port speed to 1GB and disabling auto-negotiation while an SFP is not present on an me12-10/1gb-sfp+ MDA (3HE11903AA), the physical port will remain operationally down while the transmitter laser stays on, once the 1GB SFP is inserted. To recover from this state, enable and disable auto-negotiation. To avoid this condition, only change speed and auto-negotiation settings while an SFP is present in the port. [336845-MI]
- Under certain conditions, the “Per Threshold MDA Discard Statistics” in the **show port** output does not accurately reflect the actual packet discards. [338476-MA]

13.4.2 Satellites

- When rebooting the system in model-driven mode, a conflict can arise between the configuration of Ethernet satellites and associated infrastructure and existing persistency indices files, causing an error in the boot process. Configuration of Ethernet satellites using model-driven interfaces should only be used while the system is operating in mixed mode. [336280-MA]

13.4.3 Classic CLI

- The **show**>**debug** output will output ‘router "Base"‘ without any debug configured. This is a benign CLI issue and debug is not running. [339602-MI]

13.4.4 System

- A s36-400gb-qsfpdd XMA, operating in a SR-s series router, may reset due to changes introduced in Release 19.10.R1. [330321-MA]
- On an FP4 XMA, it is possible that ports experiencing FCS errors can impact PTP performance. A workaround is to enable **rs-fec-mode** on the port where this is supported. [332619-MI]

- For IPv6, when **bof autoconfigure** is enabled, the output of the **show**>**system**> **information** command can show an invalid output for the DNS servers. As a workaround, the **show**>**router**>**management**>**autoconfigure>dhcp-client**/ **dhcp6-client**>**interface**>**management** commands can be used to show the correct DNS server addresses. [333259, 334721-MI]
- Polling the SNMP MIB “ipForwarding” object returns an incorrect value of 0 instead of 1. [334589-MI] **[NEW]**
- When a **file** command connects to a remote server over HTTPS, if the request URI specifically refers to the server's IPv6 address, and the request is made through a HTTP forward proxy, then the system will fail to verify the server's SSL certificate. [339360-MI]
- Transferring a file greater than 2 GB to the router's local storage is not completed when using Secure File Transfer Protocol (SFTP). [343084-MI]

13.4.5 NETCONF

- The combined nokia-conf/nokia-state revision dates returned in the reply to a get/filter/modules-state request do not match the actual revision dates in the combined nokia-conf/nokia-state YANG files. [342185-MA]
- Some Nokia SR OS submodules are missing in the reply to a get/filter/modules-state request. [342186-MA]

13.4.6 Model-driven Interfaces

13.4.6.1 System

- Configuration rollback is not supported in mixed and model-driven mode when configuring mirror source with LAG ports. [334871-MA]
- YANG models that only have a header and no other statements should be removed. [338502-MI]
- In rare cases, when using command authorization, a **validate** or **commit** command can cause an active CPM reset. This issue is only present in Release 19.10.R1 and 19.10.R2. This issue can only occur if local AAA profiles are being used with “match” statements that match on specific instances of objects (for example, **match** “**configure port 1/1/1**”). If AAA profile match statements do not include list key values (for example, the “1/1/1” value in the previous example) then this issue does not occur. [343002-MA]

- Saving a large configuration immediately after performing an ISSU, before the software upgrade is complete, writes an empty configuration. [343370-MA]

13.4.7 Ingress Multicast Path Management

- When multicast traffic is received by a BIER Bit Forwarding Router (BFR) and Ingress Multicast Path Management (IMPM) is enabled on the receiving card's FP, IMPM incorrectly accounts for the packet size minus 40 bytes. This causes IMPM to account for less ingress multicast traffic than is being received, which could result in packet loss instead of streams being blackholed by IMPM when the traffic on the associated paths and planes approach their maximum capacity. Packet loss can be avoided by lowering the multicast plane capacity to allow IMPM to blackhole streams when the maximum path or plane capacity is reached. [335929-MI]

13.4.8 DHCP

- In certain scenarios, when **use-gi-address scope pool** and overlapping pool names are configured, the **local-dhcp-server** may return an IP address from an incorrect subnet. [338239-MA]
- The "Remaining Potential Exp. Time" should be the sum of the configured **lease-time** and **rebind-time** values, but after a CPM High-Availability switchover the value is incorrectly increased with an additional **lease-time** value. The "Remaining LifeTime" is still correct. [339499-MI]

13.4.9 IS-IS

- The D-bit should be set for all prefixes (local and not local) advertised by the router due to propagation and/or leaking. [334881-MA]

13.4.10 BGP

- The **show router bgp routes mcast-vpn-ipv6 hunt all** command does not display Internal or Local routes. [263221-MI]

- In BGP PIC, when a BGP import policy with **install-backup-path** is used to install the backup path and the primary path goes down and comes back up, in some cases traffic does not switch back to the primary path. If then the backup path goes down, traffic will be black-holed. The work-around is to not use the policies to install the backup path. [325508-MA]
- In some interop cases of BGP Virtual Private Wire Service (VPWS) with third-party vendor routers, a third-party router may send an offset (VBO) of 0 for the VPWS Network Layer Reachability Information (NLRI). When the NLRI is withdrawn, an SR node acting as a Route Reflector (RR) propagates the withdrawal incorrectly to its clients. The standby CPM on the RR may reset, and the clients receiving withdrawal may reset the BGP session to the RR. [326598-MA]
- BGP does not advertise a unique label for routes with the same next-hop and a different **admin-tag**. [336069-MA]
- When a BGP **next-hop-resolution** policy is configured and **use-bgp-routes** is changed from enabled to disabled, it is possible that some BGP next-hops are incorrectly considered as not reachable. [338680-MA]
- An inactive aggregate route configured with the **summary-only** option can in some cases impact advertisement of the component routes. [338946-MI]
- When a CPM High-Availability switchover is completed on the node with the BMP collector, a **shutdown** must be performed, followed by a **no shutdown** under the **config**>**bmp**>**collector** context, to have the BMP collector setup TCP connections again. [339369-MA]

13.4.11 MPLS/RSVP

- If a duplicate system IP address exists in the network and MPLS RSVP LSPs are configured to use **fast-reroute facility node-protect**, it is possible for MPLS resources to get leaked. A CPM High-Availability switchover is required to recover the resources. [320821-MI]
- The following issues are present in the YANG state output for MPLS: [341366-MI]
  - "An IPv4/IPv6 operational down reason (MPLS interface and MPLS base context) is not populated if the MPLSv4/v6 (MPLS base context) and/or a MPLSv4/v6 interface (MPLS interface context) operational state becomes operational down.
  - A record label against the first egress hop for an RSVP-TE LSP is incorrectly set to a value of 4294967295.
  - "An actual route hop list and computed hop list for a localCspf LSP incorrectly shows the hop as loose with "isLoose true".

13.4.12 QoS

- A port scheduler policy that has multiple levels configured within a group, with one of those levels having no algorithmic active children, can in rare cases with a scaled number of queues, result in a line card reset. [336038-MA]

13.4.13 Services General

- Layer 2 headers are incorrectly decapsulated from an L2TPv3 tunnel when **pw-type ethernet-vlan** is used. [339388-MA]
- For circuit emulation Epipes, the system may transmit frames across the packet switched network with incorrect content in the SSRC field of the RTP header. This may cause the far-end equipment to discard these packets if it checks the SSRC value in the RTP header. Possible workarounds are: [339915-MI]
  - Configure the remote equipment to ignore the SSRC value
  - Configure the Epipe to exclude the RTP header
  - Execute the configuration **no rtp-header** command followed by the **rtp-header** command on the Epipe SAP

13.4.14 Subscriber Management

- GTP S11 hosts are removed by a hard reset of the line card on which the associated FPE is located. [318328-MI]
- If the subscriber interface addresses and/or prefixes are less than /96 (for example, 2001::/95) and the subscriber interface is in 128-bit WAN mode, GRT-leaking will fail to export routes related to host lookup. As a result, downstream traffic forwarding will fail (from network to subscriber). In order to use subscriber interface 128-bit WAN mode and GRT-leaking, addresses and/or prefixes larger than or equal to /96 are recommended. [322482-MA]
- After a double CPM High-Availability switchover on the standby SRRP side of a redundant BNG pair, subscriber routes that are normally resolved using the **redundant-interface** command can be incorrectly resolved using the local **group-interface** command. [337277-MA]
- When the **user-name-format** is either **dhcp-vendor** or **circuit-id**, then a **mac-format** of "ieee" is not supported. During an upgrade a code sets the **mac-format** value to its default value "alu". [338901-MA]

- A tools triggered Change of Authorization (CoA) message (**tools**>**perform**>**subscriber-mgmt**>**coa**) fails when a MAC address is used as key for **alc-subscr-id** and **alc-brg-id**. [342098-MI]
- PPPoE sessions connected using redundant Ethernet satellite (**esat**) ports can, upon the failover of the redundant uplink, have the keepalive timers timed out and PPPoE sessions disconnected. This can occur when keepalives are only generated locally and not by the PPPoE clients. [342122-MA]

13.4.15 VXLAN

- In an all-active network-interconnect-vxlan Ethernet Segment (ES) scenario where an R-VPLS service is configured with a VXLAN instance and EVPN-MPLS, the non-DF PE may incorrectly forward routed unknown unicast traffic to the VXLAN instance. Routed unknown unicast traffic refers to the traffic coming from the VPRN, with a resolved ARP/ND entry but unknown destination MAC in the FDB. [334472-MI]

13.4.16 IPsec

- The description of the tables tmnxIPsecTnlHistStatsTable and tmnxIPsecRUTnlHistStatsTable in TIMETRA-IPSEC-MIB is incorrect. [340770-MI]

13.4.17 NAT

- NAT Point-to-Point Tunneling Protocol Application Layer Gateway (PPTP ALG) is not working correctly in combination with **filtering address-and-port-dependent**. [339537-MA]

13.4.18 Application Assurance

- Malformed Arena of Valor game play sessions may cause the **isa-aa** to reboot. [341200-MA]

- In **isa-aa** and **isa2-aa** deployments, AA internal buffers may be leaked if the analysis of a UDP or TCP session is terminated because AA detection took longer than 15 minutes, and the first packet received after the 15 minutes is a payload packet. When all buffers are exhausted, AA can no longer provide protocol and string based application detection. [343834-MA]

13.4.19 BFD

- Switching a scaled number of MPLS LSPs from non-standby secondary paths back to the primary paths may result in a traffic loss of up to 100 ms. [334604-MA]

13.4.20 Cflowd

- Self-Generated Traffic (SGT) QoS for Cflowd applications are not marking the Cflowd packets correctly starting from Release 16.0.R4. [306019-MA]
- Enabling Cflowd unicast sampling for an ESM group-interface can result in an upstream L3 traffic drop. This applies to L2TP Access Concentrators (LAC), L2TP tunnel switches (LTS), L2-aware NAT subscriber hosts, and subscriber hosts that use a GTP uplink. WLAN-GW, vRGW L2-aware, and vRGW IPv6 firewalls are impacted as well. There is no traffic dropped in case **cflowd-parameters sampling multicast** is enabled. [340871-MA]

13.4.21 VSR

- A received packet with 64 bytes of size is incorrectly counted as an undersize packet in the port statistics. However, the packet is still processed correctly. [339235-MI]

13.5 Release 19.10.R2

13.5.1 Hardware

- For MDA types me3-200gb-cfp2-dco, x12-400g-qsfpdd, s36-400gb-qsfpdd, x6-200g-cfp2-dco, and me6-400gb-qsfpdd, provisioning a connector from **no breakout**, or **breakout** c1-100g, c4-10g, c10x10g, c1-10g, or c1-40g to **breakout** c2-100g may result in a small (less than 80ms) traffic hit and a link bounce on any connectors that are on the same MAC chip. See the **show datapath** *slot-number* command output for a list connectors and their MAC chip numbers in that slot. [335627-MA]

13.5.2 System

- The record **complete-network-interface-ing-egr** is erroneously configurable within an accounting policy. An accounting policy configured with this record cannot be applied on any object in the system. [335580-MI]
- Ingress traffic on an FP4 based MDA or XMA could be pre-classified wrongly as all-BE traffic after a soft reset (**clear card** *slot* **soft**) of the line card. This could result in unexpected traffic drop of high priority traffic in case the entire ingress FP is congested, but it will not affect normal QoS operation if there's no such congestion. [335847-MA]
- For the 4-complex XMA of the 7750 SR-1s/2s/7s/14s, it is not possible to apply a Pay-As-You-Grow software upgrade from the 3600g-dd to the 4800g hardware capacity levels. The cr3600g-cr4800g, er3600g-er4800g, and he3600g-he4800g upgrade entitlements are intended for these upgrade paths, but they are incorrectly blocked in the system. [336796-MA]
- When the **file** command connects to a remote HTTPS server through a forward proxy, if the proxy cannot establish a SSL tunnel to the server and returns a "403 Proxy Error" message, the CLI session locks up. [339445-MA]

13.5.3 MD-CLI

- If a configuration file contains blank lines at the beginning of the file, the MD-CLI **load** command displays spurious errors while loading the file. [333589-MI]

- When issuing a **load** or **rollback** command in the MD-CLI with command accounting enabled (that is, RADIUS or TACACS+), the last configuration command in the file is the only command that is logged. [337546-MI]

13.5.4 NETCONF

- When the Nokia 'combined' modules are configured to be used: [337083-MA]
    - A <get><filter><netconf-state><schemas> request returns the Nokia 'non-combined’ conf/state modules instead of the Nokia 'combined' conf/state modules.
    - Cannot acquire the "nokia-li-conf.yang" and "nokia-li-state.yang" through a <get-schema> RPC as it tries to acquire the non existing "nokia-li-conf-combined.yang" and "nokia-li-state-combined.yang" files.
- LSP-Egress statistics state information for the SR_TE LSPs and LSP-template are missing in NETCONF query results. [337718-MA]

13.5.5 Model-driven Interfaces

13.5.5.1 System

- The YANG “when” statement logic for elements that do not have a default statement in the YANG model has been changed so that the XPath expression correctly defines the constraints in cases where leafs do not exist in the configuration. For example, all “when” statements with "a != b" are changed to "not (a = b)". [333471-MI]
- In a model-driven saved configuration the **hybrid-buffer-allocation** configuration on a port that is linked to an interface with a filter configured will not boot. Remove the **hybrid-buffer-allocation** configuration, then add it back after booting. [336499-MA]
- Configuring NAT64 li-source entries (Lawful Intercept) in mixed **configuration-mode** may cause additional NAT64 li-source entries to appear in other li-sources in the classic CLI engine. If this configuration is saved to compact flash (**configure li save**), it may prevent the router from successfully executing the configuration file upon reboot. Nokia recommends not using the NAT64 li-source entries when operating in mixed **configuration-mode**. [336593-MA]

13.5.5.2 Services

- When a VRPN spoke-interface is bound to an SDP with **weighted-ecmp** enabled, switching to mixed/model-driven fails due to an audit failure. Configuring the same in mixed or model-driven mode is not allowed. [330873-MA]
- A model-driven configuration file that contains an **mc-ring** statement might fail to execute after a node reboot. [335102-MA]

13.5.5.3 Routing

- Routes with label 0 are incorrectly returned with the label-value 1036239 in the **state router bgp rib label-ipv4 local-rib** output. [335705-MA]

13.5.6 Routing

- The system incorrectly allows IPv6-in-IPv4 and IP-over-GRE-over-IPv4 tunnel termination to a local interface IP address in the CPM. [326304-MI]
- In rare cases, a FIB overflow on an FP4 line card can result in a line card reset. [333037-MA]
- When some packets in an IPv4 flow are fragmented by an egress line card, there is a small possibility that packets are transmitted out of order. This issue is only present on FP3-based cards. [333986-MA]

13.5.7 IS-IS

- In an SR-LDP stitching setup with a large number of SR-LDP stitched tunnels, the tunnels are fast to converge after a CPM High-Availability switchover. But if a second CPM switchover occurs shortly (less than 10 minutes) after the first, the tunnels could be slow to converge. [333277-MI]

- When micro-loop avoidance is running after an event and a second event occurs on the same link, the timer of the first micro-loop avoidance is not canceled. Therefore, the procedure of micro-loop avoidance is not restarted. If the second event occurs on another link, there is no issue. If the second event occurs on the same link but is handled in the same SPF as the first event, there is no issue. If the second event is the same kind as the first event (for example, an increase of the metric or deletion of the link/decrease of the metric or addition of the link), the issue is there but the micro-loop avoidance will remain the same, therefore the duration of the timer is shorter. [335912-MI]
- A node reboot may occur when the CPM runs out of memory while the **micro-loop-avoidance** command is enabled. [336631-MA]

13.5.8 OSPF

- In specific third-party vendor environments with Segment Routing enabled, the OSPF neighbor can remain in Loading state when an Extended Link LSA containing an empty Extended Link TLV is received. [337273-MA]
- The **config**>**router**>**ospf**>**seg-rtng**>**egr-stats**>**[no] node-sid** configuration incorrectly updates the next-hops for the backup-SID to use the same next-hops as the SID which it is the backup for. Therefore, it is not acting as a backup anymore. [338556-MA]

13.5.9 BGP

- It is recommended to remove the **advertise-inactive** configuration prior to modifying the **cluster-id** configuration when 6PE routes are present. Note that BGP peers will bounce when **advertise-inactive** is modified. [334446-MA]
- In scenarios where there are multiple VPN next-hops to the same IP address, it is possible that IP address is not installed in the FIB based on **selective-label-ipv4-install**. [334872-MA]
- Enabling **bgp**>**next-hop-resolution**>**use-bgp-routes** inside a VPRN while BGP routes are already present in the RIB can result in invalid routes in the RIB. As a workaround, the **next-hop-resolution**>**use-bgp-routes** option should only be enabled when the RIB is empty (for example, when **bgp** is **shutdown**). This is only an issue in VPRNs and not in the Base router instance. [335294-MA]
- When the CPMs are in the ISSU state, deleting and re-executing the Long-Lived Graceful Restart (LLGR) or **enable-notification** configurations of the peer may result in a standby CPM reset or a BGP session reset. [335370-MA]

- Overlapping BGP and aggregate routes, with BGP having a higher RTM preference, may cause the BGP neighbor to flap or the standby CPM to reset when component routes of the aggregate are updated. [336117-MI]
- Configuring **advertise-label per-prefix** with **disable-route-table-install** in the same policy might result an invalid RIB-out entry. [336420-MA]
- For aggregate routes with **summary-only** selected, component routes marked with policy **disable-route-table-install** after they have been aggregated will continue to contribute to aggregation and do not get advertised. As a workaround to get the marked route advertised, either perform a BGP **shutdown**/**no shutdown** or mark the component route with policy **disable-route-table-install** and then aggregate. [336613-MI]
- **show**>**router**>**bgp**>**optimal-route-reflection**>**bgp-nh-info location** should only be used on existing ORR locations. Using it on a non-existing location might trigger a CPM High-availability switchover. [336777-MI]
- An IES interface local route can incorrectly be used to resolve a VPN next-hop and VPN traffic, being forwarded to prefixes using such an IES interface local route to resolve the next-hop, can result in a CPM High-Availability switchover. [337000-MA]
- A backup route marked with a **disable-route-table-install** policy is not removed from the FIB table of a CSC VPRN. [337085-MI]
- When the **to neighbor** configuration statement is used in a BGP export policy, BGP routes may be incorrectly advertised when any policy or BGP configuration changes are made. [338609-MA]

13.5.10 Segment Routing

- Error messages appear on the console when deconfiguring a static SR policy (**config**>**router**>**segment-routing**>**sr-policies**>**static-policy**) after attempting to delete it using the Route Origination Module (ROM). [330311-MA]

13.5.11 EVPN

- In an EVPN OISM scenario, a multicast flow for a specific (S,G) can be shortly interrupted in the ingress PE, if the egress PE creates a (*,G) state for the same group. [334039-MA]
- EVPN SMET routes are not properly processed in an Inter-AS model B scenario. [335303-MA]

13.5.12 PIM

- With inter-AS MVPNs, the UMH PE incorrectly removes the PMSI from the outgoing interface list (OIL) if one of the source-join BGP routes are removed even if other source-join routes are present. [330723-MA]

13.5.13 QoS

- The configuration of a SAP ingress QoS policy containing a forwarding class redirection to an FP ingress queue group policer is erroneously allowed within a subscriber profile using an HS-MDA. [332930-MI]
- The system erroneously blocks users from redirecting a forwarding class policer to an FP ingress queue group in a SAP ingress QoS policy when the policy is applied under a SAP with **multipoint-shared** enabled. [332931-MI]

13.5.14 Services General

- L2TPv3 over IPsec tunnels cannot forward Ethernet frames with a destination MAC address in the range CF:00:00:00:00:00 to CF:FF:FF:FF:FF:FF. These frames are discarded. [337332-MA]
- FIPS-140-2 mode does not work together with MACsec. [346386-MA]

13.5.15 Subscriber Management

- When a data-trigger triggers mobility for an existing relay DHCPv4 and DHCPv6 lease-state, SR OS incorrectly sends out a spoofed release to the DHCP server. [334699-MA]
- The internal self generated MAC address may not be used in a RADIUS AUTH request for GTP. If the **include-radius-attribute mac-address** is set, the request will be ignored. [337774-MI]

13.5.16 VPRN

- Routes resolved to a tunneled next-hop are blocked from leaking between VPRNs. [337310-MA]

13.5.17 IPsec

- In rare cases, when using IKEv1, the ISA configured as **isa-tunnel** or **isa-tunnel-2** might reset when deleting an IPSEC-SA (phase-2) in the process of deleting the parent ISAKMP-SA (phase-1). [334095-MA]

13.5.18 Application Assurance

- Within the web-service CLI menu, if tab-complete is performed on the **category-set-id** without a classifier specified, the active CPM may reset. For example, **config**>**app-assure**>**group**>**url-filter**>**web-service**# **classifier category-set-id** <tab>. [338320-MA]

13.5.19 OAM

- When a large number of continuous SAA tests are configured, over time, some tests may be scheduled incorrectly. In rare cases, this may result in a CPM High-Availability switchover. [335154-MI]

13.6 Release 19.10.R1

13.6.1 Hardware

- The 7750 SR-a and SR-e chassis do not support the secondary management router interface being bound to the auxiliary management port of the standby CPM (A/4 or B/4 depending on which CPM is standby). When configured, this interface is kept operationally down. [278587-MI]
- When configuring a SM or PM TTI Tx string on a CFP2-DCO optics module, the actual configuration does not take effect until the **show port otu** command is entered after the configuration. [302141-MI]
- Inserting and configuring the CFP2-DCO module without an Rx signal present causes the Tx laser to toggle off and on repeatedly with log events "Coherent Optical Alarms Set/Clear (0x08) hostTx" being generated. To stop the toggling and the log events, a workaround is to **shutdown** the port until a fiber is connected. [302384-MI]

- In some cases, the 80A status LED in an APEQ-DC-2200/2800 power supply may not turn green when the PEQ is configured for 80 amperes (**input-power-mode 80**). [305641-MI]
- The Coherent Optical Port Statistics output for CFP2-based coherent optics may incorrectly show a maximum value of 40 dB for the receive OSNR upon establishing a link-up. The statistics can be cleared to reset to the actual current values. [318396-MI]
- If **ethernet autonegotiate** is disabled when using optical 1GB SFP modules on an me12-10/1gb-sfp+ MDA (3HE11903AA), the physical port remains operationally down while the transmit laser remains on. [320557-MI]
- On a CFP2 coherent optic with the connector breakout configured as **c2-100g**, with an **sm-tti** mismatch (for example, raised at only one OTU client port), the **otu-bdi** alarm raises against both far-end client ports. [322977-MI]
- On a CFP2 coherent optic with the connector breakout configured as **c2-100g**, when one client OTU port is operationally down, there are no OTU alarm indicators of the affected OTU client ports. [322978-MI]
- Enabling Cflowd on a public GRE interface over a PXC port may cause the system to incorrectly report CHASSIS #2098 tmnxEqCardQChipBufMemoryEvent errors on the IOM/XCM hosting the PXC port. [323775-MI]
- On a CFP2 coherent optic with the connector breakout configured as **c2-100g**, the Coherent Optical Port Statistics (Elapsed Seconds) are doubled. [325154-MI]
- On a CFP2-DCO optic with the OTU TTI **mismatch-reaction** configured as **squelch-rx**, the link state bounces from operational to not-operational accompanied by "Alarm Local Fault" Set/Clear toggling. With an Ethernet **hold-time** configured to two (*2*) seconds, the link remains in the correct state, but the “Alarm Local Fault" Set/Clear continues to toggle. [327145-MA]
- FCS errors on received frames on Ethernet ports may incorrectly cause the chassis #2059 tmnxEqCardPChipError alarm to be reported against the ingress forwarding complex on 7750 SR-1/1s, and against the ingress forwarding complex of the XCMs on 7750 SR-2s platforms. [327812-MA]
- With 100G coherent optics for p1-100g-tun-b, x2-100g-tun, maxp1-100gb-cfp2, x4-100g-cfp2, and me1-100gb-cfp2 MDAs, the DWDM coherent **no rx-los-reaction** behaves as though the squelch option is configured for the **rx-los-reaction** parameter. Also, with p1-100g-tun-b and x2-100g-tun, there will be a toggling **otu los** alarm. The default squelch option does behave as expected. [328333-MI]

- Soft Reset will fail on the me2-100gb-ms-qsfp28 MDA when one or more connectors are configured with a breakout of **c1-40g**, **c4-10g**, or **c4-25g** but will not fail if both connectors are configured with a breakout of **c1-100g**. To recover, a hard reset of the MDA using the **clear mda** *mda-id* command is required. [331016-MA]

13.6.2 Satellites

- An Ethernet Ring path configured on a resilient satellite port does not recover if the primary host line card is rebooted. This behavior is not observed upon a Soft Reset. [286533-MI]
- If all communications between the 7750 SR Host and a Ethernet satellite are lost, and a **client-down** delay is configured, it may take an additional 20 seconds for the satellite to declare remaining uplinks as unavailable, and then proceed to bring any remaining clients ports operationally down. [326754-MA]

13.6.3 System

- 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]
- Some Ethernet configuration are incorrectly accepted on internal ports (for example, 1/1/fm-sub) present in ISA cards. [296806-MI]
- The output of the **file dir** and **delete** commands is not consistent for both active and standby CPMs. [310531-MI, 312448-MI]
- When executing the **file dir** command for non-existent file on active and standby CPM, the error message severity should be minor. [324393-MI]
- If **cpm-http-redirect optimized-mode** is enabled (default), certain HTTP-redirect requests larger than 2048 bytes may result in an active CPM reset. [327478-MA]
- A RADIUS server specified for AAA security will use the incorrect source IP address if the active BOF IP address is changed without a CPM reset. As a workaround, the RADIUS server configuration can be removed and added using the **config**>**system**>**security**>**radius server** command. [328827-MI]
- If a 7750 SR-1s or an XCM card in a 7750 SR-2s or SR-7s resets as a result of a software or hardware defect, a limited amount of debug information is logged in the **admin tech-support** file about this reset. [334603-MA]

13.6.4 MD-CLI

- When switching from **configuration-mode classic** to **configuration-mode model-driven** while an **isis overload timeout** is configured, the overload setting incorrectly does not expire after timeout. In **configuration-mode model-driven**, it is recommended to use **tools perform router isis overload** *seconds*. [329803-MI]
- Modifying the **mbs** and **cbs** through **qos** sap-ingress and sap-egress **overrides queue** incorrectly accepts values in kilobytes, while the value is always processed as bytes. The **?** help command also incorrectly indicates that the entered values should be entered in kilobytes. [329941-MI]

13.6.5 Model-driven Interfaces

13.6.5.1 Services

- In a service where **provider-tunnel** and either **bgp-evpn**, **bgp-ad**, or **bgp-vpls** are enabled, configuring a second BGP protocol (either **bgp-evpn**, **bgp-ad**, or **bgp-vpls**) fails in the model-driven interfaces and causes an audit failure in management interface configuration mixed **configuration-mode**. The **provider-tunnel** needs to be enabled after all desired BGP protocols are enabled. [327383-MI]
- MD-CLI validation will incorrectly pass when the same SAP is configured for both MSTP (with a **managed-vlan-list**) and L2PT termination in a VPLS service. Commit fails with an error that indicates the conflict between the managed VLANs and L2PT. [330352-MI]

13.6.5.2 OAM

- When attempting to read the OAM-PM state model, no values will be returned for IP and MPLS. [330402-MA]

13.6.5.3 Routing

- Redundant **allowed-peer-as** configurations for dynamic BGP sessions, such as an AS number and a range with minimum and maximum AS number set to the same value (**configure router bgp group** *"bgp-group-1"* **dynamic-neighbor match prefix** *0.0.0.0/0* **allowed-peer-as** *["***64496..64496***" "***64496***"]*), are accepted in the MD-CLI but should not be used to prevent inconsistency when operating in **management-interface configuration-mode mixed** or **model-driven**. [320041-MI]
- The **clear**, **show**, and **tools dynamic-services** related commands are not supported in the MD-CLI. If executed they can, in certain cases, result in a CPM High-Availability switchover. [324612-MA]

13.6.6 APS

- A classic 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 operation, even if one or more member links still remain present during the course of the rollback. [121066-MI]

13.6.7 IPv6

- Enabling the **advertise-tunnel-link** command (forwarding adjacency) 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. Furthermore, IPv6 packets which are received may be blackholed at the router enabling forwarding adjacency. 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]

13.6.8 OSPF

- In OSPFv2, when a local router receives a subnet (A) from a neighboring router through an OSPF Router LSA that can be locally resolved by subnet (B) in the same area on the local router, the local router marks the route as directly resolved instead of the correct next hop. This happens if both subnets and the link between routers are advertised in the same area. In addition, the router LSA from the neighboring router must list subnet (A) as one of the stubs after the P2P link between routers. The workaround is to advertise subnet (A) as a Type 5 external LSA from the originating router or change the links to type broadcast.

  In OSPFv3, when a local router receives a subnet (A) from a neighboring remote router through an OSPFv3 intra-area prefix LSA that can be locally resolved by subnet (B) in the same area on the local router, the local router marks the route as directly resolved instead of the correct next hop. This happens if the link between both routers and both subnets are advertised in the same area as intra-area routes. The workaround is to advertise subnet (A) as a Type 5 external LSA from the originating router. [330908-MA]

13.6.9 IS-IS

- The event tmnxIsisMaxSeqExceedAttempt IS-IS #2032 does not provide enough information on what is happening with the protocol. It should indicate that the IS-IS instance will be temporarily **shutdown**. [311679-MI]
- An IS-IS Segment Routing tunnel will be deleted and re-added if the SID flag is modified. For example, if the original SID-flag combination is 12 [NnP], and after a network event it changes to 12 [RNnp], the tunnel will be deleted and then re-added, resulting in a small impact to the traffic. [320534-MI]
- If IS-IS is performing a TLV check on ELC & RLDC codepoints which are not yet defined, a router-capability TLV 242 may be discarded when an LSP is received with the bfd-discr sub-TLV (0x14). [334029-MA]

13.6.10 BGP

- When more then 200 communities are present on a BGP route, the **show router bgp routes hunt** command can become unresponsive. [318153-MI]
- When the **bgp next-hop-resolution use-bgp-routes** option is enabled, the next-hop resolving process incorrectly does not consider any subscriber management routes. [318967-MI]

- AS_CONFED_SET is incorrectly advertised for non-confed-members EBGP peers. [323978-MI]
- For Inter-AS VPRN option B, IGP routes may be incorrectly not used to resolve the next-hop if a default static to route (**static-route-entry 0.0.0.0/0**) is present. Also, a BGP MED value can have an incorrect value if the export policy uses **metric set igp**. [328930-MA]
- With **advertise-inactive** enabled, if a policy is added that rejects a prefix (IGP to BGP), the router incorrectly does not withdraw the advertised BGP routes. Note that the BGP peers bounce when the state of **advertise-inactive** is changed. [329833-MA]

13.6.11 BGP-EVPN

- If a **spoke-sdp** is configured within a CsC VPRN, all traffic destined to the prefixes using the **spoke-sdp** as a next-hop from the CsC-CE will be dropped. A workaround is to use **auto-bind-tunnel**. [326200-MA]

13.6.12 Segment Routing

- Segment Routing (SR) label operations could be synchronized incorrectly after a standby CPM reset. As a result, unexpected traffic loss may occur during a subsequent CPM High-Availability switchover for traffic using SR tunnels. [318714-MI]

13.6.13 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. 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]
- In the **tools**>**dump**>**router**>**segment-routing**>**tunnel** output, the out-label for a stitched tunnel (for example, SR-LDP stitching) will be '0' in Release 19.5.R1 or earlier instead of '-'. The ‘0’ label does not mean an IPv4 explicit-null label and the correct out-label is the label of the LDP tunnel displayed under Interface/ Tunnel-ID. The label of the LDP tunnel can be displayed via several **show** commands, such as **show router tunnel-table**. [309319-MI]

- The router incorrectly replies with an ROUTE_REFRESH ACK when it receives an RSVP PATHERR or RESVERR containing the RSVP ROUTE_REFRESH bundle message with ACK_Desired, even though **refresh-reduction** is not configured on the corresponding local RSVP interface. [321780-MA]

13.6.14 PIM

- With **auto-rp-discovery** enabled under the PIM instance of a VPRN participating in an MVPN, if the node receives an PIM ASSERT PDU on the I-PMSI for an Auto-RP message with source as *0.0.0.0*, then the active CPM will reset. A CPM High-Availability switchover will occur on a redundant system, but if the same packet comes in again before the standby comes up, the new active CPM will also reset, causing the node to reboot. Auto-RP messages should have the source address as the mapping agent's or the candidate RP’s loopback IP address, instead of *0.0.0.0*. Such a packet could be generated spuriously by an errant third-party router. Two or more PEs receiving such a packet on the CE interfaces may pass it on the I-PMSI tunnel triggering PIM ASSERT procedure on the participating PEs. These PIM ASSERT PDUs contain the errant source address *0.0.0.0* which causes the CPM to reset. [319328-MA]
- If an interface address is removed while it is still present in the **pim** context, and the address is re-used for another interface, then enabling PIM on the new interface causes the active CPM to reset. To avoid the reset, the older interface should be removed from the **pim** context before adding the new interface. [319947-MA]

13.6.15 QoS

- When reporting reserved CBS buffer changes, the red and amber reserved CBS alarms for FP ingress network buffer pools reference the MDA (ObjType=mda and Owner=x/y) instead of the FP. [317734-MI]
- In certain cases, the output of **show port** *port-id* **detail** may display an incorrect large value for the port access available bandwidth instead of zero. This is only a display issue. [320987-MI]
- The **configure qos sap-egress ip-criteria entry** command is supported except with the configuration of **port-redirect-group-queue true** within the **action** statement. [322312-MI]

13.6.16 Services

- When a Unicast DHCPv6 packet (such as RelayFwd) enters a VPRN and a route-lookup points to the Global Routing Table (GRT) which again points to a group interface IP address of yet another VPRN, the following log messages maybe generated. The messages are harmless and can be ignored. [329602-MI]

```
“xxxxxxxxxxxxxxxxxxxxx LOGGER #2002 Base A:TIP:UNUSUAL_ERROR
"Slot A: tipIp6Input: UNICAST DHCP RECEIVED WITH GRT, BASE to 9"”
```

- Removing a LAG from under the **redundancy**>**multi-chassis**>**mc-lag** configuration context, where the LAG is being used in PBB B-VPLS, causes issues with MAC learning in both B-VPLS and associated I-VPLS instances on the node where **mc-lag** is in standby state. [330157-MA]

13.6.17 Subscriber Management

- For RADIUS-based credit control, when credit is exhausted or depleted, the system should send a RADIUS authentication message. However, the RADIUS message is incorrectly not generated in case the **authentication-policy** is applied via a **local-user-db**. [222924-MI]
- The **show redundancy multi-chassis mc-ring peer** *ip-address* **ring** *sync-tag* **ring-node** command incorrectly displays an extra entry when a 32-character length ring node name is used. [322611-MI]
- Moving populated subscribers between HS-MDA ports (ports on the HS-MDAv2 or IOM4-e-HS) and non-HS-MDA ports can result in a "sbmGetParentActSub sbmActSub <sub-id> invalid subAppId" log event generated on the active CPM. In very rare events, it can also result in a CPM High-Availability switchover or node reboot. A workaround is to clear the active subscribers before switching them between the two types of ports. [324714-MA]
- Clear text is not acceptable on NGE-configured interfaces. Traffic carrying only an "MPLS service label", but no "NGE Encryption label" is accepted, even if an inbound NGE key-group is configured on this group interface. [325297-MI]
- In dual-homed setups using Multi-Chassis Synchronization (MCS), DHCP hosts that do not renew/rebind and only send two or more discoveries when the lease expires may trigger a temporary "subMgmtIpoe lost sync with peer" alarm. [325386-MI]
- Processing Diameter messages at a high rate may cause higher-than-expected "Subscriber Mgmt" CPU capacity usage. [325485-MA]

- When receiving a modify bearer request for GTP S11 sessions, the message is incorrectly dropped if there is no 'Sender F-TEID for Control Plane' IE and the source UDP port does not equal 2123. [326211-MI]
- The setup rate for IPv6 data-triggered hosts with the same MAC address can be slow on a static SAP configured under a subscriber interface with **wan-mode** set to 128-bit or with **delegated-prefix-length** configured smaller than 64 bits. [326793-MI]
- In very rare cases, a data-triggered host with Gx Usage Monitoring enabled can get in a state without a valid Diameter session. [330872-MA]
- From Release 15.0.R4 onwards, when no SLAAC address information is returned from RADIUS or LUDB, no unsolicited route-advertisement is sent after a DHCPv4 host establishment and **ipoe-linking gratuitous-rtr-adv** is configured. [332079-MA]
- An IPv6 prefix configured on a **subscriber-interface wan-mode mode128** before another IPv6 host prefix that has a length between 32 and 95 bits is configured and then removed after the IPv6 host prefix was added may lead to high CPU usage and the bouncing of the IPv6 hosts that are instantiated under the IPv6 host prefix. This is not an issue for IPv6 host prefixes with lengths of 96 bits or longer. [334390-MA]
- In case of ESM-over-GTP, if a GTP session reconnects and the RADIUS server is not responding, or slow to respond, benign LOG events can be raised, such as: "SUBMGR:UNUSUAL_ERROR Slot A: sbmEiGetAddr: pParent == NULL" or "IPOE_SESSION:UNUSUAL_ERROR Slot A: ipoeSessionGetEsmInfo: pSession == NULL" for "sbmGtpSessionReconnect". [334855-MI]

13.6.18 IPsec

- When an IPsec tunnel phase-2 (IPsec SA) is being re-keyed, **traffic-forward** counters may be incorrect in that polling interval and the next polling interval. [314911-MI]
- The **isa** *mda* [**saved-key**] option is incorrectly missing from the **clear ipsec tunnel** command. [324501-MI]
- An IPv4 flow comprised of fragmented and non-fragmented packets, sent to an ISA (configured as **isa-tunnel** or **isa2-tunnel**) on the private SAP, may be received on the public SAP in a different order after encryption. [334254-MA]

13.6.19 VXLAN

- Packets from a VTEP received on a **restrict-protected-src discard-frame** enabled VXLAN instance are incorrectly discarded if the source MAC address in the packet matches a protected FDB entry associated with a different VTEP in the same VXLAN instance. [319571-MA]

13.6.20 NAT

- 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]
- When deterministic NAT mappings are configured on the standby node of a dual-homing NAT redundant system, the corresponding inside IP address is incorrectly populated in the RTM/FIB table after executing a **nat-group** *nat-group-id* **no shutdown** command. [333870-MA]

13.6.21 Cflowd

- Toggling reach ability of a VPRN **cflowd collector** can cause the source-address to not be modified to the IP address specified by **vprn source-address application cflowd**. [328051-MI]
- Cflowd packets that are sent to in-band collectors in the Base router instance with **inband-collector-export-only** configured are not sent with the correct IP Header length. [336088-MA]

13.6.22 OAM

- **oam lsp-trace** for a BGP labeled-unicast prefix fails when the underlying tunnel is RSVP and the last-hop router for the RSVP tunnel is configured to send **implicit-null-label**. [319667-MI]
- After 7000 GTP pings, no further OAM pings can be issued. [326161-MA]

13.6.23 WLAN-GW

- IOM or MDA redundancy for a WLAN-GW group does not recover dual-stack or single-stack IPv6 UEs. [323048-MA]

13.6.24 VSR

- The **show**>**system**>**cpu** command shows incorrect CPU usage when VSR is running on a hyper-threaded host. [325056-MI]

13.7 Release 19.7.R2

13.7.1 System

- When configuring **card power-save** for a line card that supports Soft Reset, the card is only partially reset and full power-saving is not realized for that card. A workaround to achieve optimal power-savings for a card is to use any of the following options: [328296-MI]
    - Perform a hard reset of the card with **clear card** *slot*
    - Pull out and re-insert the card
    - Reboot the chassis

13.7.2 NETCONF

- Using NETCONF <get> RPC to acquire the state data of an MDA/XMA that is equipped but not provisioned, does not return the equipped-level for that MDA/XMA. [326429-MI]

13.7.3 Model-driven Interfaces

13.7.3.1 IPsec

- Discrepancy issues exist with accepted IPv6 addresses between the classic CLI and the MD-CLI. Global unicast IPv6 addresses are only accepted for the following fields in the MD-CLI, whereas all IPv6 addresses are accepted for the classic CLI. [327939]

```
ipsec-tunnel/dest-ip
ipsec-tunnel/key-exchange/dynamic/id/ipv6
ipsec/security-policy/entry/local-ipv6
ipsec/security-policy/entry/remote-ipv6
ip-tunnel/dest-ip
```

13.7.4 BGP

- With **advertise-inactive** enabled, if a policy is added with a **prefix-list** that rejects a prefix (IGP to BGP), the router incorrectly does not withdraw the advertised BGP routes. [326959-MI]
- When BGP convergence tracking is enabled (by configuring the **min-wait-to-advertise** *seconds* to a value greater than 0), the convergence process is not tracked after the **max-wait-to-advertise** time limit is reached, and the state goes into timeout. As a result, the state stays in timeout after all connected BGP peers have sent an End-of-RIB (EOR) message. [327216-MA]
- The **show router bgp convergence** output field for the first session established time of a peer incorrectly changes with the **clear router bgp neighbor all** command. [327243-MA]
- The **show router bgp convergence** output field for the last session established time for a neighbor is incorrectly updated beyond the **min-wait-to-advertise** timer expiry and also with the **clear router bgp neighbor all** command. [327261-MA]

13.7.5 LDP

- **monitor router ldp fec-egress-stats** can display incorrect outputs for Collect Stats, Admin State, and Accounting Plcy. [327610-MI]

13.7.6 VSR

- Ingress hierarchical policing may incorrectly drop traffic below the policer rate. This is likely to occur on systems with 10 or more worker tasks, a policer with a rate of 50 Gb/s, and an offered rate above 10 Gb/s. [325059-MA]

13.8 Release 19.7.R1

13.8.1 Hardware

- On a 7750 SR-12 equipped with PEM-3s with an external un-split AC rectifier shelf, a rectifier fail alarm is only raised if two rectifiers from the same side of the power shelf experience a failure. This may result in the failure of a single rectifier on either side of the power shelf going unreported. [259116-MI]
- On a 7750 SR-7/12 with CPM5 installed, it is possible for the power supply model to be displayed as unknown in the **show chassis detail** output. This can occur if there is a CPM switchover and the SFM, that the newly active CPM is seated in, is not operationally up. [293066-MI]
- In a system equipped with apeq-dc-4275 PEQs, if any of the APEQs are running on a single feed, the system will not report an initial half-power state. However, if the APEQs start with both valid input feeds, then it will report if one of the feeds become unavailable. [318028-MA]
- With a CFP2-based coherent optics in an me3-200gb-cfp2-dco MDA, the **dwdm coherent rx-los-reaction** provisioned as **squelch** does not take effect and does not take the link down. The **otu sm-tti mismatch-reaction squelch-rx** can be provisioned as an alternative. [321070-MI]
- With a 100G coherent DWDM optics module, an SNMP query of the provisioned wavelength and channel may not update its value until a CLI query is performed when the port is set as **admin shutdown**, or it may take up to 60 seconds to update when the port is set to **no shutdown**. The actual provisioning does take place, only the display update is affected. [320900-MI]
- Starting with Releases 15.1.R1 and 16.0.R1, the mechanism to monitor certain hardware errors on the CPM5, CPM-S, CPM-X16, and CPM-X20 is not functioning correctly in case the associated error occurs. As a result, the following log events are not generated on the CPM: CHASSIS #2059 tmnxEqCardPChipError, CHASSIS #2063 tmnxEqCardPChipMemoryEvent, and CHASSIS #2076 tmnxEqCardPChipCamEvent, and no CPM High-Availability switchover is triggered. [323141-MA]

13.8.2 Satellites

- Setting a connector on a satellite to **shutdown** (or admin-state to **disable**) does not change the operational state of the ports on that connector. [322687-MI]

13.8.3 System

- A standby CPM may remain in syncing state after a CPM High-Availability switchover if there are a large number of ports that have LLDP enabled. When in this state, a "Slot 2: lldpQueueAsyncMsg: Queuing LLDP Activity message failed" log event will be generated and to recover from it, the line card with slot number mentioned in the log event should be reset. If an ISSU (which requires a CPM switchover) is planned on a node with a lot of LLDP enabled ports, it is recommended to remove LLDP prior to the ISSU and add it back after. [319728-MI]
- When provisioning a breakout connector, the system does not check if enough egress queues are available for the new ports. If the number of egress queues is depleted on any forwarding complex, the system might fail to execute such configuration upon reboot with the error message: "MINOR: PMGR #1023 Not enough resources - egress queue". [322385-MA]
- When configuring TACACS+ node management inside a VPRN service, the system uses the **priv-lvl-map priv-lvl user-profile-name** from the **system security aaa remote-servers tacplus** instead of the one configured under the service. If there is no **priv-lvl-map priv-lvl user-profile-name** configured under the **system security aaa remote-servers tacplus** context, the "default" profile is received. [326168-MI]
- CPM cards are incorrectly not displayed in the output of **show system switch-fabric**. [326783-MI]

13.8.4 Model-driven Interfaces

13.8.4.1 System

- A **cflowd collector export-filter interface-list** configuration is missing from the MD-CLI when switching from **configuration-mode classic** to **configuration-mode model-driven**. [327371-MI]

- MD-CLI commands in CRON, EHS, and VSD scripts will incorrectly fail to execute (rejected by command authorization) when no **cli-user** is configured in the **configure system security cli-script authorization** [**cron** | **event-handler** | **vsd**] **cli-user** *user-name* context. A workaround is to create a user with full access to all commands needed in the scripts, and assign that user as the **cli-user**. [337228-MA]

13.8.4.2 NETCONF

- The nokia-notifications.yang file is missing from the "/YANG/" directory of the shipped NOKIA SR OS YANG modules. [323981-MA]

13.8.5 Routing

- In gRPC-based RIB/FIB API, a CPM switchover does not preserve the currently remaining time on the RIB-API purge timer. The timer is reset and starts counting down again from its configured value. [309523-MI]

13.8.6 IP/RTM

- When FRR is activated at the ingress LER of an SR-TE tunnel, CPM-originated packets over the SR-TE LSP are dropped. [264579-MA]
- A static-route over an SR tunnel with ECMP next-hops (including BGP routes resolved over such static-route) has incorrect information for the second ECMP next-hop onwards. The detailed next-hop information is erroneously copied. [321859-MA]

13.8.7 BGP

- An **aggregate** route with an **indirect** *ip-address* that resolves to an IPv4-over-IPv6 route is not supported. [319616-MA]
- Routes are not advertised if the next-hop applied in BGP export-policy is incompatible with the family type. For example, if an IPv6 next-hop is specified for an IPv4, label-IPv4, or VPN-IPv4 route then the route is only advertised towards RFC 5549-capable peers, and not towards the other non-RFC 5549 peers. [321887-MI]

13.8.8 MPLS

- For SR-TE LSPs, MPLS incorrectly reports LSP status as down to PCC after a successful path computation (PC) update MBB. As a result, PCC and MPLS become out of sync. [322811-MA]

13.8.9 Subscriber Management

- An **admin rollback revert** fails when the configuration contains a **subscriber-mgmt shcv-policy** *name* **trigger** *trigger-type* statement. [313932-MI]

13.8.10 Mirroring

- On a 7950 XRS chassis, if the destination mirror ports are on slot 9 or higher, then any other XCM with 200G or 400G XMAs, whose ports are used as mirror sources, may reset. [326122-MI]

13.8.11 NAT

- The active-active NAT redundancy mechanism might fail when a CPM High-Availability switchover occurs, followed by an ISA failure. [320188-MA]

13.8.12 Application Assurance

- The ISA (configured as **isa-aa** or **isa2-aa**) may reboot when presented with SMB traffic with unexpected segmentation. [324329-MA]

13.9 Release 19.5.R2

13.9.1 Hardware

- There is a small possibility that a line card will reset when taking an **admin tech-support**. FP4-based cards are not affected. [315881-MA]

13.9.2 Satellites

- Collecting an **admin tech-support** file can result in a line card reset if an Ethernet satellite (**esat**) port is assigned to a multi-service site (MSS) that has two uplinks to different IOMs/IMMs. [321344-MI]

13.9.3 MD-CLI

- On an interface **flavor unnumbered-mpls-tp**, configuring a **lag-per-link-hash** is not supported, even if it is the default **lag-per-link-hash** configuration. [321190-MI]
- In B-VPLS services, the low-order two bytes of the operational B-MAC are reserved if **use-mclag-bmac-lsb** is configured; however, in the MD-CLI there is no validation check if the operational B-MAC low-order two bytes and the MC-LAG **source-bmac-lsb** do not match. In mixed-mode, this missing validation may cause the classic CLI to MD-CLI configuration file conversion to fail. [322098-MA]

13.9.4 Ingress Multicast Path Management

- When the **bandwidth-policy** configured under **fp** *fp-number* **ingress mcast-path-management** is changed with a different number of secondary paths, the queue scheduling priority of the paths that changed between primary and secondary is not updated to correspond to the new path assignment. As a workaround, after the **bandwidth-policy** is changed, remove it from **fp** *fp-number* **ingress mcast-path-management** and then re-add it. [321835-MI]

13.9.5 OSPF

- In certain scenarios using OSPF Segment Routing with multiple areas, it is possible that **loopfree-alternates** may not be correctly established. To prevent this issue, a loopback interface with a node-SID should be configured in each area. This issue is not applicable to IS-IS Segment Routing. [318699-MI]

13.9.6 BGP

- The Peer AS field of the **show router bgp neighbor** output displays a value of "0" for Dynamic BGP neighbor. [318561-MI]
- It is possible that some IPv4 label-BGP routes are not included in a multipath set of routes in the VPRN route-table despite configuring the **ignore-nh-metric** or **multi-path unequal-cos**t. This typically occurs when some of the BGP next-hops resolve over a local route while other BGP next-hops resolve over a route from another protocol. [320274-MA]

13.9.7 BGP-EVPN

- The **vpls**>**bgp-evpn**>**vxlan**>**auto-disc-route-advertisement** command is not supported on a service where any VXLAN instance is associated with a **network-interconnect-vxlan** Ethernet Segment. This combination is incorrectly allowed by the system and may result in an active CPM reset. [321176-MA]

13.9.8 QoS

- In certain cases, the output of **show qos scheduler-hierarchy** may display an incorrect large value for the port scheduler max rate when it is overridden to **max**. This is only a display issue. [320787-MI]
- When two or more ports which are on the same forwarding complex are added to a LAG, traffic over these ports may not be correctly classified by the **action policer** in the SAP egress QoS policy. [321813-MI]

13.9.9 Subscriber Management

- In a highly-scaled L2oGRE setup, packet drops can be observed upon an SDP **shutdown**/**no shutdown**. [322039-MA]
- The Acct-Session-Time attribute within subscriber service RADIUS accounting messages is no longer updated when one of the hosts of a dual-stack PPPoE or IPoE session is deleted in the system. [322924-MA]

13.9.10 NAT

- When using the **tools dump nat sessions** command in conjunction with a filter on the field **inside-ip**, the filtering of the output is incorrect. [319904-MI]

13.9.11 Application Assurance

- Under unexpected traffic conditions, DNS tunnel resources may be leaked. This may result in a reboot of the ISA. [317828-MA]
- An **app-filter** expression entry that uses "\\" to represent a single "\" in a uniform resource identifier (URI) is an accepted expression, but it is rejected by the ISA (configured as **isa-aa** or **isa2-aa**). This results in a configuration that will never provide a URI match. A workaround is to use the percent encoding of %x5c to represent the "\" in the URI expression. [320405-MI]
- When a CPM High-Availability switchover occurs, either the last or second last configured AA DNS-IP-cache domain entry will not be downloaded to the ISA. This will result in no IP address entries in the DNS-IP-cache for the missing domain. A workaround is to configure two additional dummy entries to ensure all of the valid entries are downloaded correctly. The dummy entries must be the last entries configured, so it is recommended to use entries that start with "z" to ensure that they are configured last if the node reboots. [321703-MA]

13.10 Release 19.5.R1

13.10.1 Hardware

- The **show system switch-fabric** command does not display the correct value for fabric capacity for the following license levels. The issue is specific to the **show** output and capacity alarms, not affecting the actual value. [295568-MI]
    - 1.6T and 3.6T **cr**, **er**, and **he** licensing levels on 7750 SR-7s/14s
    - 1.2T and 1.6T **cr**, **er**, and **he** licensing levels on 7950 XRS-20/20e
    - 400G **cr**, **er**, and **he** licensing levels on 7750 SR-12e with IOM5-e
- Beginning with Release 16.0.R5, the 1PPS output port on CPM5 of the 7750 SR-7/12/12e platforms is not functioning as expected. While it still presents a pulse every second, it is not connected to the internal PTP time-base and cannot be used to evaluate time accuracy on the CPM. [306704-MI]
- When inserting a CFP8 optics module, a series of log events may be generated indicating multiple transitions between operational and non-operational states. [317909-MI]

13.10.2 Satellites

- An Ethernet satellite (**esat**) port that is assigned to a multi-service site (MSS) with a **scheduler-policy** attached, can trigger an IOM/IMM reset in certain scenarios when secondary **esat** ports are added and removed again from the **port-map**. [305766-MA]
- When SyncE is enabled in Ethernet satellite configuration and 100G link is used as uplink port, the satellite fails to boot up upon soft reboot (**admin satellite reboot**). This issue is seen on the 7210 SAS-Sx 10/100 64SFP+4CFP4 variant only. It is not applicable to the 7210 SAS-Sx 10/100 64SFP+4QSFP28 variant. A workaround is to disable SyncE before the soft reboot and then re-enable SyncE after the satellite has returned to service. If the satellite is soft-rebooted with SyncE enabled, then a power-cycle of the satellite is required to recover. Refer to the *7210 SAS Software Release Notes* to identify the specific SAS OS release needed for the resolution. [308806-MA]
- In rare cases, for traffic crossing an Ethernet satellite, the packet payload of an IP fragment, which is not the initial fragment, may be corrupted if all of the following conditions exist:
    - the Ethernet satellite is enabled for transparent clock functionality

  - IPv4 packet with IP protocol value identifies the packet as UDP
  - bytes 3 and 4 after the IP header is equal to 319 or 320
  - lower 4-bits of the 9th byte after IP header is equal to 0, 1, 2, or 3
  - lower 4-bits of the 10th byte after IP header is equal to 2

  Refer to the *7210 SAS Software Release Notes* to identify the specific SAS OS release needed for the resolution. [309868-MI]

13.10.3 System

- After a CPM High-Availability switchover, a log event may not have the correct slot information and will be displayed as: "Status of Card ? changed administrative state: inService, operational state: outOfService". [314985-MI]

13.10.4 NETCONF

- There are discrepancies between the YANG modules advertised in: “/netconf-state schemas” and “/modules-state” versus the distributed YANG files. For example, nokia-conf-aa-common and nokia-conf-aa-group-part-policy are advertised but not distributed. [303449-MA]

13.10.5 Model-driven Interfaces

13.10.5.1 System

- A request for the state information of internal ports (specifically in ISAs) has the unicast, multicast, and broadcast packets statistics missing in its output. These are categorized as Ethernet statistics. [298320-MI]
- Configuration of a port connector breakout in model-driven interfaces may require multiple transactions (multiple candidates and commits) for certain operations. For example, a change of a breakout requires transactions to delete all configurations that are dependent on the breakout (for example, ports, SAPs, LAGs, and router interfaces), a transaction to delete the **breakout**, a transaction to configure the new **breakout**, then transactions to rebuild the attributes removed in the first step.

- Configuration of a card or an MDA in model-driven interfaces may require multiple transactions (multiple candidates and commits) for certain operations. For example, changing an MDA type requires transactions to delete all configurations that are dependent on the MDA (for example, ports, SAPs, LAGs, and router interfaces), a transaction to delete the **mda-type**, a transaction to configure the new **mda-type**, and then transactions to rebuild the attributes removed in the first step.

13.10.5.2 Services

- The L2TPv3 configuration incorrectly allows a local interface address to be configured as an L2TPv3 local-address as well. This type of configuration should be avoided as it may cause the system to become unstable. [314657-MA]

13.10.6 Routing

- The QoS forwarding class of incoming IP multicast packets at the ingress PE with DSCP bits set (to BE, for example) is not carried over to the MPLS EXP bits of the outgoing BIER data packet's header. [318038-MA]

13.10.7 BGP

- Rapidly performing the sequence of **clear**, **shutdown**, and **no shutdown** commands on a BGP peer may lead to an incorrect AIGP value for a locally-imported BGP route. [312884-MI]
- The advertisement of a static-route for prefix 0.0.0.0/0 cannot be manipulated with a **vsi-export** policy when matching on a **prefix-list**. A workaround is to assign a tag to the static route and match on the tag value in the **vsi-export** policy. [316008-MI]
- Receiving a BGP route that covers its own next-hop may cause high CPU usage in the BGP task on the standby CPM. [316223-MI]
- Executing in classic CLI edit-cfg mode a candidate load with configuration line **bgp neighbor local-address** *ip-address*, displays a benign "Referencing non-existing object '*ip-address*'" warning. [316762-MI]
- CLI configuration of BGP **graceful-restart long-lived advertise-stale-to-all-neighbors without-no-export** in classic CLI edit-cfg mode may result in an active CPM reset. [316995-MI]

13.10.8 LDP

- The LDP PW FEC 128 optional interface parameter sub-TLV VLAN tag is incorrectly not reset in case there is a withdraw received for the label mapping. [317805-MI]

13.10.9 QoS

- 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, 289026-MA]
- The "Egress Policers" and "Egress User Queues" rows in the **tools dump resource-usage card** *slot-num* **fp** *fp-number* output have been moved to after the "Dynamic Policer Stats (in use by Egress)" row in the same output. [308953-MI]
- In rare cases, deleting a secondary shaper on an HS-MDAv2 can fail when ESM MSAP subscriber is being created at the same time with the **inter-dest-id** referencing that secondary shaper. [314627-MI]
- The creation of SAP ingress and egress QoS policies with names beginning with anything other than letters (either lowercase or uppercase) are erroneously rejected. Exceptions include entries beginning with a number or an underscore followed by alphabetic characters, and a numeric name with leading + sign. [319032-MI]
- If an Ethernet satellite (**esat**) port has a SAP with shared-queuing configured, deleting a secondary uplink can result in a benign error message: "SVCMGR:sapGetPortMemberIngHwResources NumSharedQueueObjs[14] can't go negative!". [319100-MI]

13.10.10 Services General

- IES/VPRN spoke interfaces are not supported when the spoke-SDP binding resolves to a hierarchical tunnel which consists of an LDP or BGP Label Unicast (BGP-LU) tunnel which resolves to an IGP-shortcut using an SR-TE LSP. [278837-MA]

- IES/VPRN spoke interface, R-VPLS interface, and EVPN services are not supported when the spoke-SDP binding or EVPN destination resolves to any of the following hierarchical tunnels: [278834-MI]
  - LDP FEC which itself resolves to an IGP-shortcut using an SR-TE LSP
  - an SR-ISIS tunnel, SR-OSPF tunnel, or SR-TE LSP which itself resolves to an IGP-shortcut using an RSVP-TE LSP

13.10.11 Subscriber Management

- A **traceroute** from a retail subscriber host, destined to a retail VPRN BGP-VPN route, will fail to resolve the first hop and displays "1 0.0.0.0 * * *". [312829-MI]
- For ESM-over-GTP, reporting of ECGI or TAI changes in Gx is inconsistent when only one of TAI or ECGI is received from MME. [314469-MI]
- On the S11 interface, the following sequence of GTP messages may result in a reset of the standby CPM:
  - Create a session for an IMSI on APN1
  - Create a session for the same IMSI on APN2
  - Modify Bearer Request and Response for IMSI on APN1
  - Delete Session Request and Response for IMSI on APN1
  - Modify Bearer Request for IMSI on APN2

  In such a case, the following event can be logged: "redTraceFailedUpdate: Upd M[93]:GtpApi D[6]:GtpApiDownlinkLearned SET". [319687-MI]

13.10.12 VPRN

- When NGE is enabled on a VPRN, and **grt-leak** is configured to provide IP reachability from the VPRN to the GRT (for example, for remote system management through the VPRN), packets coming from the GRT to the VPRN are unencrypted. As a result, the packets are discarded at the receiving end of the VPRN. [307256-MA]

13.10.13 VXLAN

- Deleting an Epipe with a **vxlan-src-vtep** configured can result in a failure to create new services afterwards. As a workaround, the **vxlan-src-vtep** configuration should be set to the default value before the Epipe service is deleted. [299325-MI]

13.10.14 NAT

- Creating PPPoE-client BRGs with **port-forwarding-range** value configured as the maximum range (65535) in a **nat-policy** may cause ISA instability. A workaround is to configure any values less than the maximum range. [316547-MA]

13.10.15 OAM

- For an IPv4 header override, which does not specify source or destination address values, the values that are not specified are set to 0.0.0.0 after the second CPM High-Availability switchover. [305401-MI]

13.10.16 VSR

- In some rare cases, after a fatal error, the system fails to reset automatically. While in this state, the system drops to a minimum number of cores and operates in a degraded state. [309748-MA]

14 Known Issues

Following are specific technical issues that exist in Release 19.10.R6 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.

14.1 Hardware

- The optics modules details displayed in the output of the **show port detail** 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]
- When the 7950 is operating as an XRS-40 and some ports are to be configured for synchronous Ethernet, then the BITS ports of the Extension chassis must be cabled to a synchronous source. If there is no need for synchronous Ethernet ports, then there is no need to cable the BITS ports. However, currently a benign critical alarm will be raised against the BITS ports on the extension chassis if they are not cabled. There are two workarounds. [192096-MI]
  - Set the **event-control** for *tmnxEqSyncIfTimingHoldover* and *tmnxEqSyncIfTimingHoldoverClear* to “suppress”
  - Cable the BITS ports on the extension chassis using one of the two configurations shown in the installation guide and configure the sync-if-timing bits output to source internal-clock

- 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]
- The 100G **dwdm coherent dispersion** configuration is dependent upon the **dwdm coherent mode** being configured as **manual**. The dispersion configuration must be configured prior to the **dwdm coherent mode** configuration, or the dispersion configuration will not take effect. [318397-MI]
- When present in me1-100gb-cfp2 MDA, the CFP2-DCO optic can show a Transceiver Digital Diagnostic Monitoring (DDM) low warning threshold crossing event log for its supply voltage. [334759-MI]

- When using a c10-10g breakout connector, and one of those ports is used as a reference into the central frequency clock, then events occurring on that port can impact the other nine ports of the module. Configuring and unconfiguring the port as a reference, the reference port entering link down, or the reference port entering link up causes a short interruption on the other nine ports (link-down followed by link-up). It is recommended to not use ports on a c10-10g as a reference for the central clock. [345502-MA]
- The addition of thousands of CPM filter entries on the 7750 SR-1 or 7750 SR-1s can result in a node reset. [349240-MI] **[NEW]**
- The output of **show**>**chassis**>**power-management** incorrectly displays provisioned satellites. [349518-MI] **[NEW]**

14.2 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]
- A configuration **rollback revert** in classic CLI that changes the primary link of a resilient satellite client port which has SAPs or router interfaces configured on it does not complete successfully on the first attempt. A second rollback is necessary to complete successfully. [277319-MI]
- After a **clear** of an MDA or a line card hosting the primary uplink for a resilient satellite client port, a larger-than-expected traffic loss may occur during the revertive switchback from the secondary to the primary uplink. [282226-MI]
- A configuration rollback from a VRRPv3-configured setup with satellites to a plain configuration without VRRPv3 may fail. [282580-MI]
- A configuration rollback (in the classic CLI) from a QSFP satellite with connector ports configured to one without the QSFP satellite will fail. [302554-MI]
- SAPs bound to an Ethernet satellite client port cannot be associated with a multi-service sites with **assignment card** type. [305943-MA]

- When a port-template is created in model-driven mode without having its ports fully defined, and then a switch is performed to classic mode, the port-template will enter an invalid state and should not be applied to the satellite from classic mode. A workaround is to remove the port-template from any satellites where it is applied, re-create the port-template, and then re-apply it. [336018-MI]
- For an Ethernet satellite client port with a resilient **port-map** setting, it is possible to add the Ethernet satellite client port to a LAG when at least one of its uplinks has a valid **port-topology** statement. Once the Ethernet satellite client is a member of the LAG, it is possible to remove the resiliency by removing the secondary uplink. If the secondary uplink is the one with the valid **port-topology**, then the Ethernet satellite port becomes unhosted while it is part of the LAG. This situation can lead to a High-Availability CPM switchover when the Ethernet satellite port is removed from the LAG. A workaround is to remove the Ethernet satellite port from LAG prior to changing its **port-map**, then to add it back to the LAG. [346435-MA]

14.3 Classic 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 classic CLI **admin rollback revert** operation fails toward a configuration with **maximum-routes** *value* or **maximum-ipv6-routes** *value* defined. [252516-MI]

14.4 System

- 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]

- CPU-protection policies are not supported at the IES/VPRN tunnel interface SAP-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]
- For IMMs and MDAs that support IEEE 1588 Port-Based Timestamping (PBT), after an ISSU, the CPM may expect the port to execute port-level timestamping of the PTP frames, although the IMM or MDA is not running the up-to-date firmware that supports this feature. This may result in corrupted correction fields in the PTP messages. This only impacts PTP ports (Ethernet encapsulation) and not PTP peers (IP encapsulation). To resolve this issue, the firmware should be upgraded using the **clear card** or **clear mda** command. The firmware version can be verified with the **show** command of the assembly (for example, **show mda** *mda-id* **detail**). This issue affects Release 14.0.R4 and later. [228493-MI]
- 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 SFMs 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]
- 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]
- Setting a user-defined CPM queue's **mbs** and **cbs** value to a maximum of 131MB is incorrectly allowed. [271667-MI]
- When the link on a 100G CFP is down, the reason is incorrectly displayed as "channelNotConfigured". There should not be any reason displayed. [276288, 290479-MI]
- If a **sync-if-timing** source-port that is selected for use stops receiving SSM packets, and the reference received quality level goes into a failed state, the **wait-to-restore** period does not take effect as long as source-port does not go physically down and the receipt of SSM packets resumes on the source-port. [304144-MI]
- When a port is in the signal-degrade or signal-fail condition and the tmnxEqPortEtherCrcAlarm log event has been configured for repetition, the log-event is raised repeatedly but will stop repeating if there is a CPM High-Availability switchover. [305435-MI]
- BITS output squelching (when enabled) is not occurring on the standby CPM/ CCM BITS output port when the input line reference has a quality level lower than the configured **ql-minimum** and output of CPM clock identifies the source of BITS output port (**bits output source internal-clock**). [305783-MI]
- BITS output squelching (when enabled) is not occurring on the active and standby CPM/CCM BITS output port when the BITS input reference has a quality level lower than the configured **ql-minimum** and output of CPM clock identifies the source of BITS output port (**bits output source internal-clock**). [306052-MI]
- The following tables do not invoke configuration traps. As a result, the "User Last Modified" field in the output of the **show system information** is not updated. [306243-MI]
  - tmnxPortTable
  - tmnxQosPoolAppTable
  - tBgpInstanceTable
  - vRtrIfTable
  - vRtrLdpIfTable
  - vRtrOAM
  - vRtrRipInstanceTable

- While the system is processing a configuration file using the **save**, **exec** or **load** commands, the custom hash should not be configured or re-configured using the **admin system security hash-control custom-hash** command as it may lead to incorrect hashes. While the configuration file is being processed, the custom hash is not able to modify any hash within the configuration file with the new custom hash value, resulting in the old hash being loaded into the system. This old hash is not decrypted with the new custom hash key and the previous custom hash key. [313812-MI]
- If MACsec is configured on the me2-100gb-ms-qsfp28, performing a Soft Reset causes MACsec traffic to be lost in one direction. A workaround a **shutdown**/**no shutdown** of the MACsec sub-port after the Soft Reset, to deconfigure MACsec from the me2-100gb-ms-qsfp28 prior to the Soft Reset, or to perform an IOM hard reset. [324094-MA]
- If the SyncE/1588 port (**synce**) is enabled as one of the input references to the SETS (**sync-if-timing** or **central-frequency-clock**) and the BITS output port is enabled, then the BITS output port will not output a valid signal. There is no work around to make BITS output work when **synce** is enabled.

  In addition, if BITS settings like the interface-type are modified while **synce** is enabled as SETS reference, it will cause the **synce** reference to go into a LOS condition. A workaround to restore **synce** functionality is to disable **synce** and then enable it again. If **synce** was enabled and later disabled and BITS was then enabled, then the BITS input reference may be in a faulty condition like LOS or AIS and the BITS output port may present and invalid signal. A workaround to restore BTS functionality is to toggle the **bits interface-type** (for example, if the interface type was **e1 pcm31crc**, it can be set to **ds1 esf** and then back to its original value **e1 pcm31crc**). This workaround is applicable to both the BITS In and BITS Out when **synce** is no longer enabled but was previously enabled. [332287, 334370-MA]
- Under **config**>**li**>**li-source** or **li**>**li-source**, SAPs can be configured before the SAP exists under any service, known as loose reference. In this case, the port encapsulation type must not be changed. If the port encapsulation type is changed, then it must be followed by the deleting or re-creating all of the associated LI source SAPs. Any **admin save** in this scenario must follow a **li-save**. [334621-MA]
- On the 7950 XRS, under **configure system sync-if-timing**, when **synce** is enabled while the **bits interface-type** is set to **ds1 esf** and the BITS input port is receiving an E1 signal, then the **synce** reference will fail to qualify. This occurs even when the BITS input is disabled. When using the **synce** input do not connect any signals to the BITS input port. [334948-MA]

- The SR OS currently allows the SSM **code-type** to be programmed as SONET on one of the **synce** ports (for example, A/3) and SDH on the other port (for example, B/3). This causes a conflict with the **configure system sync-if-timing synce ql-override** selection and can result in a failure on processing the configuration file on a reboot. Both ports should be set to the same SSM **code-type.** [336709-MA]
- The system currently allows to change the **ssm code-type** of the SyncE ports (for example, A/3) after a **config**>**system**>**sync-if-timing**>**synce**>**ql-selection** was programmed. This causes a conflict between the new **ssm code-type** and the existing **ql-override** and can result in a failure to process the configuration file on a reboot. The **ssm code-type** of the SyncE ports (for example, A/3) should be changed only if **config**>**system**>**sync-if-timing**>**synce**>**ql-selection** is not set. [336720-MA]
- After a double High-Availability CPM switchover, the available licenses are not displayed in the output of **show**>**system**>**license**>**available-licenses**. A workaround to correct this display issue is to use the **admin**>**system**>**license**>**validate** command to restore the information. [350047-MA] **[NEW]**
- Incorrect CPM card slot names (A swapped with B) may be included in the following log event: “The standby CPM A has a different memory size than the active CPM B”. [353692-MI] **[NEW]**

14.5 Model-driven Interfaces

14.5.1 Common Issues

This section lists items that apply to all MD interfaces.

14.5.1.1 System

- Validation does not find invalid configurations for the **peq**, **power-shelf**, and **power-module-type**. However, an error is returned when an invalid configuration is committed. On state side, no error message is displayed when state information is requested from invalid chassis, **power-supply**, **peq**, **power-shelf**, and **power-module-type**. [302137-MI]
- Persistent system indices for chassis MAC address pools in model-driven management interface configuration mode are not supported, and may change after a system reboot.

- In model-driven configuration mode, only the latest saved configuration file is synchronized to the standby CPM. This means the previous configurations are not available for rollback after a CPM switchover occurs. As a workaround, the previous configuration files can be manually copied to the active CPM, or the path to the compact flash of the previously active CPM could be used in a **load full-replace** command. [320418-MI]
- A read request to “state/chassis/power-supply” returns no power-supply statistics on platforms that do not support a configurable power supply. Retrieval of this information using SNMP is not affected. [330582-MI]
- The AAA “administrative” profile was updated with entry 50 to match **admin**>**system**>**security** in Release 16.0.R6. If the router was configured to use model-driven management interface configuration mode in Release 16.0.R5 or older, and booted or loaded with an old configuration using **load full-replace** in Release 16.0.R6 or newer, the profile entry will be missing and the administrator will not be able to access the **admin**>**system**>**security** commands in the classic CLI. The entry must be added manually to give the administrator access to the **admin**>**system**>**security** commands. [334859-MI]
- In Release 19.10.R1 the output of the **admin show configuration** command executed in the LI configuration region context, or if the LI configuration region was entered with **edit-config**, does not display the running configuration. The command must be executed at the operational root context, and after executing the **quit-config** command if the explicit configuration workflow was entered, to display the running configuration. [335161-MI]

14.5.1.2 QoS

- Modifying the **md-auto-id qos-policy-id-range** command is not supported when a network QoS policy is applied under the ingress or egress of an Epipe spoke-SDP, VPLS spoke-/mesh-SDP, IES interface spoke-SDP, VPRN interface spoke-SDP, Base router, or a VPRN network interface. [309408-MA]

14.5.1.3 Routing

- The next-hop of black-holed BGP routes is not correctly displayed in the state routing table. [299083-MI]

- For the 7950 XRS platforms, the **multicast-redirection** command in the **configure router policy-options policy-statement entry action** context is configurable and accepted by the CLI, but is not applicable to the XRS platforms. It does not provide any functionality. The **multicast-redirection** feature is only applicable in Enhanced Subscriber Management (ESM) which is not supported on XRS platforms. Do not configure this command in this context. If already configured, Nokia recommends removing the command from the policy. [317689-MI]

14.5.1.4 IMPM

- The **configure multicast-management chassis-level per-mcast-plane-capacity total-capacity** element in MD interfaces incorrectly accepts a value of 16500 starting from Release 16.0.R2 until Release 19.10.R1. When using **configuration-mode mixed**, this value can be saved to the classic startup configuration file which causes an error, halting the processing of the configuration file on a reboot. [345589-MI]

14.5.2 MD-CLI

This section lists items that apply only to the MD-CLI.

- The **compare** command output for user-ordered lists is incorrect and the desired state is not achieved by pasting the **compare** command output back into the MD-CLI. [350843-MA] **[NEW]**

14.6 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-XP 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-XP 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 or in NETCONF (by configuring "member *vpi*" or "member *vpi*/0"), but this action is blocked in the CLI. [243704-MI]

14.7 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 an existing **trap-target**. A workaround is to remove the **trap-target**, then reconfigure the **trap-target**. [313286-MI]

14.8 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]

14.9 PTP

- When using the 10x10G breakout connector, some ports may introduce a time error in the PTP messages transiting those ports. The time error can be up to 100 ns. [297299-MI] **[NEW]**

14.10 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 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]

14.11 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]

14.12 Routing

- The **show router bgp routes** *prefix* **vpn-ipv4** [**hunt** | **detail**] CLI command incorrectly does not display prefixes with a route-distinguisher (RD) equal to 0:0. [275409-MI]

14.13 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]
- An **admin-tag-policy** can not be used as policy variable. [291590-MI]

14.14 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]
- In certain scenarios, the **local-dhcp-server** may return different values than the configured ones for the *lease-time*, *lease-renew-time*, and *lease-rebind-time*. [326902-MI]
- A DHCP server configured in **failover access-driven** mode that has a high scaled number of leases with very short lease-times can, upon a CPM High-Availability switchover, drop DHCP messages because of hitting an overload state. In this case, the following events may be logged: "Server name could not allocate IP address to client. Reason: Too many messages queued" or "*server-name* dropped a packet from client. Reason: MCS busy with previous update". [339496-MA]

14.15 IP/RTM

- Transit IPv6 packets forwarded over IPv4 IGP-shortcuts and Forwarding Adjacencies by default have the TTL of their outer MPLS label set to 255 instead of inheriting the IPv6-header's hop-limit value. Applying the **configure router ttl-propagate label-route-transit** command will enable the TTL propagation.

  CPM-originated packets forwarded over IPv4 IGP-shortcuts and Forwarding Adjacencies by default have the TTL propagated (set same as IPv6-header's hop-limit) to the outer MPLS label. Applying the **configure router mpls no shortcut-local-ttl-propagate** command will disable the TTL propagation.

  TTL propagation from the IPv6 packet header to the IPv6 Explicit Null Label (label value 2) is always performed and cannot be changed. [254050-MI]
- InBroadcastPkts, HCInBroadcastPkts, and OutDiscards counters are not incremented when broadcast or multicast traffic is sent over an interface not configured with an IPv4 address. [320260-MI]
- The key calculation for AES128-CMAC-96 and HMAC-SHA1-96 for the TCP Authentication Option (TCP-AO) may not calculate a key consistent with RFC 5926. This issue may affect interoperability with other implementations of the TCP-AO. [344165-MI]

14.16 IS-IS

- When used in combination with ECMP, the **show router isis lfa-coverage** command may provide incorrect results. [142527-MI]
- No TI-LFA backup path is found when the node SID programmed to the P-node is part of another IS-IS level compared to the SR tunnel toward that P-node. [339303-MA]

14.17 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]
- The number of extended communities displayed in the **show router bgp routes flow-ipv4** and **show router bgp routes flow-ipv6** commands is limited to ten. [262669-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]
- When routes are stored in the route-table, they have a single metric that applies to all next-hops. With unequal-cost ECMP features (for example, **bgp best-path-selection ignore-nh-metric** and **bgp multi-path unequal-cost**) the different next-hops of a route can have different metrics but only one of these values is displayed as the metric of the route. This metric comes from the last next-hop that was installed. [320593-MI]
- With multiple ECMP paths available for BGP address-family **vpn-ipv4** or **vpn-ipv6**, and if the metric of some of the those ECMP routes are dynamically changed to make them better than other routes, the ECMP paths selected as best-path may be incorrect. Also using the **configure service vprn** *id* **bgp best-path-selection ignore-nh-metric** or **configure service vprn** *id* **ecmp-unequal-cost** command during this time can result in an incorrect selection of multi-paths. A workaround is to perform a **shutdown** followed by a **no shutdown** on the BGP or VPRN instance. [321776-MI]
- Re-executing a policy configuration with **global-variables** of the type **number** used in main-policy will trigger redundant BGP updates. [328140-MI]
- The peer AS number in the BGP syslog message may be shown as 0 if the peer is either in a **shutdown** state or is not configured. [334543-MI]
- An **sr-policy** with a top label as **bgp-sr** is not supported. [337255-MI]
- Enabling the selective download feature in combination with a **bgp dynamic-neighbor** does not function correctly. [339226-MI]

- In cases where the number of active LSPs is greater than the configured ECMP value, changing the ECMP value can result in the packets not being classified as per CBF. The packets are distributed equally across all the available ECMP paths. [341600-MA]
- When a route policy change is applied to a peer that is down, no actions from that policy are applied as expected. However, if one of those actions is set to **disable-route-table-install**, the policy action is actually applied to the stale routes (LLGR/GR) from the peer after an HA switchover occurs. [341829-MI]
- Self-Generated Traffic (SGT) to a 6PE route with a next-hop over a 3107 BGP tunnel, which resolves over a LDP-over-RSVP tunnel, is dropped. [348942-MA]
- The **show**>**router**>**bgp**>**neighbor** *ip-address* **detail** command shows “Client Reflect” as disabled when it is actually enabled and vice versa. [349400-MI] **[NEW]**
- When ORF is negotiated between BGP peers and Route-Target Constraints (RTC) is later configured on only one of the peers, a route is not imported into a VPRN if its **vrf-target** is changed on the peer configured with RTC. The workaround is to ensure that RTC is configured on both peers. [350188-MI] **[NEW]**
- In some interop cases of BGP VPLS with third-party vendor routers, a third-party router may send an offset (VBO) of 0 for the VPLS NLRI. If the NLRI is withdrawn with a VBO of 0, the message is parsed incorrectly and the NLRI remains in the RIB-IN. [350394-MA] **[NEW]**
- When a classic CLI **rollback revert** operation is performed, the BFD session for any BFD-enabled BGP neighbor is cleared. [352247-MI] **[NEW]**
- Using **advertise-inactive** causes advertised labels for 3107 BGP labeled routes to change if there are changes in the corresponding RIB-IN routes. The issue is not seen if all corresponding RIB-IN routes have path IDs attached. [352918-MI] **[NEW]**

14.18 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]
- In a service with two BGP instances, if multiple MAC/IP routes are received in the same instance for the same MAC, different IPs, but same priority (based on the EVPN selection criteria), the first route to come in is selected by EVPN. In this case, only the first route is redistributed to the other BGP instance, while all routes for the same MAC and different IPs should be redistributed. [288293-MI]
- Associating a **network-interconnect-vxlan** Ethernet-Segment (ES) to a VXLAN instance in a service without **bgp-evpn mpls bgp 2** is not supported, but is incorrectly allowed by the system. In this case, the **vxlan-instance** in the service is taken to MhStandby and the MAC routes are purged. If such service is an R-VPLS with **ip-route-advertisement**, when the **network-interconnect-vxlan** ES association is removed from the service, the system fails to re-add the IP-prefix-derived MAC routes to the R-VPLS service. The workaround is to remove the **vprn**>**int**>**vpls**>**evpn-tunnel** and re-add the command. [302464-MI]
- An ASBR or next-hop-self RR router supporting EVPN address family re-advertises EVPN routes upon a CPM switchover. [324560-MI]

14.19 Segment Routing

- CPM-originated packets, including OAM packets, are not supported when the underlying IPv6 SR-TE LSP has its first hop matching a local adjacency with link-local address only and with no global unicast address. Specifically, SDP keep-alives time out and an SDP using the IPv6 SR-TE LSP, as well as any services using this SDP, goes operationally down. [345442-MA]
- SR-TE LSPs may experience a small traffic hit when reverting from a non-standby secondary to the primary path. [338702-MI]

14.20 MPLS/RSVP

- 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]
- If the **one-to-one** option is configured in **fast-reroute**, when one detour is merged into another detour, there is currently no check that the avoid node of the merging detour is not in the path taken by the detour into which it is merging. If a detour is erroneously merged into another detour which transits through the avoid node of the merging detour, then the result is a loop. If a network event activates a detour along this loop, then the result is stale RSVP sessions on the nodes within the loop but the LSP itself will go down. [287785-MI]

- Modifying the **class-forwarding forwarding-set policy** on an LSP is accepted in the CLI, but a delete/add sequence is recommended to ensure the new policy becomes effective immediately. [331843-MI]

14.21 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 **[aes-128-cmac-96]** and **[hmac-sha-1-96]** are among the supported TCP encryption algorithms for LDP and BGP sessions. 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]
- With weighted ECMP, if an ECMP tunnel next-hop resource in FIB or dataplane (IOM) is exhausted locally, LDP is not notified (overload feature) of this failure and might cause a traffic outage for the affected FECs. [272367-MI]
- An LDP IPv4 session may not come up when interoperating with a third-party LDP implementation and the **legacy-ipv4-lsr-interop** CLI command is enabled on the SR OS LSR. Nokia recommends disabling any LDP IPv6 configuration on the SR OS LSR interfaces. [298234-MI]

14.22 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]

14.23 QoS

- When a **sap-ingress** QoS policy containing a **mac-criteria** statement is applied to a routed VPLS SAP, the system incorrectly allows an **action** statement to be configured in the **mac-criteria entry** context. [322626-MI]
- When the ingress or egress QoS policy of an SLA profile is changed from a QoS policy that is used only once (in one SLA profile) on any FP in the system to a QoS policy that is not yet used on any FP in the system and when the **ip-criteria** resources on an FP are nearly depleted, the resource management on the system can go into an inconsistent state.

This issue occurs for **ip-criteria** resources of QoS policies configured in SLA profiles and applied for subscriber hosts (see the ingress or egress QoS entries of the **tools**>**dump**>**resource-usage**>**card**>**fp** command). It occurs both for **ip-criteria** and **ipv6-criteria**, with or without the use of an **ip-prefix-list** or **ipv6-prefix-list**. It occurs only for the **ip-criteria** resources; it does not occur for other QoS resources such as queues or policers. [349356-MI] **[NEW]**

14.24 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]
- Addition or removal of prefix-exclude entries to or from **ip-prefix-list** or **ipv6-prefix-list** might lead to CAM resource management failures when the line card CAM is full or almost full. [313665-MI]

14.25 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]
- For circuit emulation Epipes, the system may transmit frames across the packet switched network with incorrect contents of the SSRC field of the RTP header. This may cause far-end equipment to discard these packets if it checks the SSRC value in the RTP header. Possible workarounds are: [327692-MI]
  - Configure the remote equipment to ignore the SSRC value
  - Configure the Epipe to exclude the RTP header
  - Execute the configuration **no rtp-header** command followed by the **rtp-header** command on the Epipe SAP
- SDP **weighted-ecmp** is not supported when SDP has **sr-te-lsp** configured. [331531-MA]
  - **weighted-ecmp** will not be honored with Layer 2/3 traffic flows

  - CPM-generated ping/OSPF adjacency fails for VPRN/IES spoke-SDP bindings
- Even though a manual SDP with an IPv6 **far-end** address can be configured under a VPRN service, the indirect next-hop BGP route resolution does not work and traffic is not forwarded correctly. As a workaround for the VPRN to use the manually configured SDP, **auto-bind-tunnel** can be configured. [336365-MI]
- Traffic load balancing can be unequal in a BGP EVPN VPLS service with auto-bind tunnels using equal cost multi-path load balancing, resolved with BGP over either an RSVP or SR-TE LSP. [336864-MI]

14.26 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]
- 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]
- For subscriber interfaces using 128-bit WAN mode, configuring a static IPv6 WAN host sharing a /64 prefix is not supported. [292480-MI]
- Downlink buffering and paging for idle ESM-over-GTP connections does not work on 7750 SR-s platforms. [312538-MA]
- LAC subscriber host status may not change to the "Not Fwrding" state after a LAG is **shutdown**, resulting in potential traffic loss. [322806-MA]
- During the subscriber-host instantiation phase, various parameters (like DNS servers) for hosts are gathered from single or multiple sources. Normally the parameter sources are in order of priority: LUDB, RADIUS, DHCP SERVER, and so on. However, if the Local Address Assignment (LAA) is used it always overrules all other sources. [346012-MI]
- In case a **hold-time**>**down ip|ipv6** *value* timer is applied on a subscriber interface, ESM proxy scenarios in a single BNG topology can result in blackholing traffic for a period of maximum *value* seconds. [347035-MI]
- After disabling **arp-populate** in a working IPoE scenario, managed ARP entries may reappear in the ARP table after a DHCP renew cycle. [350186-MI] **[NEW]**

14.27 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]
- In scaled PBB B-VPLS scenarios using MRP, a node can become slow to respond to some SNMP get requests. [341641-MI]

14.28 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]

14.29 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]

14.30 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 very rare cases, the processing of an (S,G) entry on a **video-sap** may cause the active CPM to reset and result in a CPM High-Availability switchover. [291538-MI]

14.31 NAT

- 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. [200170-MA]
- 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]
- Traffic counters in the **nat inside downstream-ip-filter** are zero after a CPM switchover. [235940-MA]
- When configurations and logs are removed by a rollback, some traps contain "<NULL>" instead of the VPRN-id. [335187-MA]

14.32 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]

14.33 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]

14.34 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]
- Setting the **config**>**saa**>**test**>**trap-gen**>**probe-fail-threshold** (tmnxOamPingCtlTrapProbeFailureFilter) and the **config**>**saa**>**test**>**trap-gen**>**test-fail-threshold** (tmnxOamPingCtlTrapTestFailureFilter) to 0 during a ping style SAA test may result in the ProbeFailedV3 trap, TestFailedV3 trap, or both being generated, even if there were no failures. [208234-MI]
- For option B inter-AS and intra-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 VPRNs 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 VPRN. [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-tll-propogate** and **shortcut-transit-ttl-propagate** for SR-TE LSP shortcuts are unsupported. [262159-MA]
- 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. [296367-MI]

- Although broadcast and multicast IP addresses can be configured for IPv4 and IPv6 headers and overrides of the Egress Finder feature, such addresses are not supported by the Egress Finder. [305912-MI]
- When an SR-TE LSP, having an underlying SR LFA backup switches to the LFA backup, OAM **sr-te lsp-trace** over the LFA backup next-hop will time out. [324810-MA]
- Unusual stats may be reported by a TWAMP-light session when duplicate packets are received at the reflector. [337884-MI]

14.35 VSR

- On the VSR, egress weighted round-robin scheduling does not take effect when parented to a virtual scheduler. The egress scheduling is affected in the following cases: [297301-MA]
  - when egress queues are parenting a virtual scheduler directly
  - when egress queues are parenting a virtual scheduler and the virtual scheduler is port-parented to a port-scheduler level
- If using a virtIO port on KVM and CentOS 7.4 or later, it may be necessary to enable allmulticast or promiscuous mode on the macvtap associated with the virtIO port in order to receive multicast packets such as OSPF Hello messages. [330236-MI]
- In PCI pass-through mode, the port Input Error and Output Error counters do not increment (they remain at 0) in the presence of errors. [345667-MI]
- Egress statistics are not supported on the VSR. [349155-MA] **[NEW]**

15 Change History for Release 19.x Release Notes

The following table lists significant documentation changes to the *SR OS 19.x Software Release Notes*.

Table 36 **Change History**

Part number	Date of Issue	Reason for Issue and Changes to Documentation
3HE 15407 0010 TQZZA 01	June 2020	Sixth 19.10 Release Notes.
3HE 15407 0009 TQZZA 01	April 2020	Fifth 19.10 Release Notes.
3HE 15407 0008 TQZZA 01	March 2020	Fourth 19.10 Release Notes.
3HE 15407 0007 TQZZA 01	February 2020	Third 19.10 Release Notes.
3HE 15407 0006 TQZZA 01	December 2019	Second 19.10 Release Notes.
3HE 15407 0005 TQZZA 01	October 2019	First 19.10 Release Notes.
3HE 15407 0004 TQZZA 01	September 2019	Second 19.7 Release Notes.
3HE 15407 0003 TQZZA 01	July 2019	First 19.7 Release Notes.
3HE 15407 0002 TQZZA 01	June 2019	Second 19.5 Release Notes.
3HE 15407 0001 TQZZA 01	May 2019	First 19.5 Release Notes.

Customer Document and Product Support

Customer Documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation Feedback

Customer Documentation Feedback