



7210 Service Access Switch Release 9.0 Rev. 04

7210 SAS OS Software Release Notes 9.0R4

3HE11497AAADTQZZA V9.0.R4

Issue: 02

April 2017

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or trade names of their respective owners.

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

© 2016-2017 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Release Notice	5
1.1	About this Document	5
1.2	Release 9.0.R4 Documentation Set	5
2	Release 7210 SAS supported hardware	7
3	Changed or Deprecated Commands.....	9
3.1	Release 9.0.R1.....	9
4	Software Upgrade Notes	29
4.1	Upgrade to 9.0R1 or later releases	29
4.2	Upgrade to 8.0R1 or later releases	30
4.3	Upgrade to 7.0R4 or later releases	31
4.4	Upgrade to 7.0R1 or later releases	31
4.5	Upgrade to 6.0R6 or later releases	32
4.6	Upgrade from 6.0R4 or prior releases	32
4.7	Upgrade from 3.0 releases	33
4.8	Upgrade from 4.0 or 5.0 releases	33
4.9	Upgrade from 5.0R1 or 5.0R2 releases.....	34
4.10	Upgrade to 5.0R1 or later releases from prior releases	35
4.11	Upgrade to 4.0R4 or later release from prior releases	37
4.12	Upgrade to 1.1R7 or later releases from prior releases	38
4.13	Upgrade to 1.1R6 or later releases from prior releases	38
5	Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, Mxp, and Sx (standalone mode) devices.....	41
5.1	Standard Software Upgrade Procedure	41
6	Software Upgrade Procedures for 7210 SAS-R	45
6.1	Standard Software Upgrade Procedure	45
7	Resolved Issues.....	49
7.1	Resolved in R9.0R4.....	49
7.2	Resolved in R9.0R3.....	50
7.3	Resolved in R9.0R2.....	50
7.4	Resolved in R9.0R1.....	51
8	New Features	53
8.1	Release 9.0R4.....	53
8.2	Release 9.0R3.....	68
8.3	Release 9.0R2.....	74
8.4	Release 9.0R1.....	74

9	Enhancements	85
9.1	Release 9.0R4.....	85
9.2	Release 9.0R3.....	85
9.3	Release 9.0R2.....	85
9.4	Release 9.0R1.....	85
 10	 Known Limitations.....	 89
10.1	ACLs.....	89
10.2	CLI	90
10.3	CES	91
10.4	DHCP	91
10.5	IGMP Snooping	91
10.6	IP	92
10.7	LAG	93
10.8	Management.....	93
10.9	MPLS.....	94
10.10	Mirror	94
10.11	OAM	95
10.12	Routing	97
10.13	QoS	98
10.14	Security.....	100
10.15	Timing.....	101
10.16	Services.....	102
10.17	Statistics	104
10.18	STP.....	106
10.19	System.....	106
 11	 Known Issues.....	 109
11.1	ACLs.....	109
11.2	CLI	110
11.3	CES	110
11.4	IGMP Snooping	110
11.5	IP	111
11.6	LAG	112
11.7	Management.....	112
11.8	MPLS.....	113
11.9	Mirror	114
11.10	OAM	114
11.11	QoS	116
11.12	Timing.....	119
11.13	Services.....	121
11.14	Statistics/Accounting	122
11.15	STP.....	123
11.16	System.....	123
11.17	Hardware	125

1 Release Notice

1.1 About this Document

This document provides an overview of the Service Access Switch Operating System (7210 SAS OS) in Release 9.0R4 for 7210 SAS- D, E, K, M, X, T, Mxp, Sx, R6 and R12 Platforms.

1.2 Release 9.0.R4 Documentation Set

The 7210 SAS OS Release 9.0.R4 documentation set consists of Release Notes and the user guides.

Table 1 Release 9.0.R4 Documentation Set

Document title	Platform	Part number
7210 SAS M, T, X, R6, R12, Mxp, Sx OS MPLS Guide	M, T, X, R6, R12, Mxp, Sx	3HE11496AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OS Routing Protocols Guide	M, T, X, R6, R12, Mxp, Sx	3HE11491AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OS Router Configuration Guide	M, T, X, R6, R12, Mxp, Sx	3HE11495AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OS System Management Guide	M, T, X, R6, R12, Mxp, Sx	3HE11489AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OS Basic System Configuration Guide	M, T, X, R6, R12, Mxp, Sx	3HE11487AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OAM and Diagnostics Guide	M, T, X, R6, R12, Mxp, Sx	3HE11493AAAATQZZA
7210 SAS M, T, X, R6, R12, Mxp, Sx OS Interface Configuration Guide	M, T, X, R6, R12, Mxp, Sx	3HE11485AAAATQZZA

Table 1 Release 9.0.R4 Documentation Set (Continued)

Document title	Platform	Part number
7210 SAS D, E, K OS Routing Protocols Guide	D, E, K	3HE11490AAAATQZZA
7210 SAS D, E, K OS Router Configuration Guide	D, E, K	3HE11494AAAATQZZA
7210 SAS D, E, K OS System Management Guide	D, E, K	3HE11488AAAATQZZA
7210 SAS D, E, K OS Basic System Configuration Guide	D, E, K	3HE11486AAAATQZZA
7210 SAS D, E, K OS OAM and Diagnostics Guide	D, E, K	3HE11492AAAATQZZA
7210 SAS D, E, K OS Interface Configuration Guide	D, E, K	3HE11484AAAATQZZA
7210 SAS D, E, K OS Services Guide	D, E, K	3HE11478AAAATQZZA
7210 SAS M, T, Mxp, Sx OS Services Guide	M, T, Mxp, Sx	3HE11479AAAATQZZA
7210 SAS X, R6, R12 OS Services Guide	X, R6, R12	3HE11480AAAATQZZA
7210 SAS D, E, K OS Quality of Service Guide	D, E, K	3HE10379AAACTQZZA
7210 SAS M, T, Mxp, Sx OS Quality of Service Guide	M, T, Mxp, Sx	3HE11482AAAATQZZA
7210 SAS X, R6, R12 OS Quality of Service Guide	X, R6, R12	3HE10381AAACTQZZA

2 Release 7210 SAS supported hardware

NOTE:

Refer the SR 7x50 documentation to know which 7210 release is operating in satellite mode and which 7x50 host release needs to be used for different variants of the satellite.

- 7210 SAS-S 48F 4SFP+ and 7210 SAS-S 24F 4SFP+ platforms in standalone mode are supported from 7210 SAS 9.0R4 release.
7210 SAS-K 2F4T6C (12-port IP/MPLS device) Non-ETR platform is supported from 9.0R3 release. Referred as 7210 SAS-K 2F4T6C in this documentation.
- Non-POE 7210 SAS-Sx 48T4SFP+ and 7210 SAS-Sx 24T4SFP+ platforms are supported from 9.0R1 Release.
- 7210 SAS-Sx 48T4SFP+ PoE and 7210 SAS-Sx 24T4SFP+ PoE platforms are supported from 8.0R8 Release.
- 7210 SAS-Sx platform is supported from 8.0R6 release.
- Following Part numbers of IMM-SAS-R-b are supported from 7210 SAS 8.0R6 release.
 - IMM-SAS-R-b 10SFP-1SFP+ (P/N:3HE09152AARD)
 - IMM-SAS-R-b 2SFP+ (P/N: 3HE09153AARD)
 - IMM-SAS-R-b 4SFP+ (P/N: 3HE09154AARD)
 - IMM-SAS-R-b 11cSFP (P/N: 3HE09155AARD)
 - IMM-SAS-R-b 16TX (P/N: 3HE09156AARC)
- From 8.0R3 release, 7210 SAS-R12, 7210 SAS-Mxp and 7210 SAS-T supports 8GB compact flash (3HE04708AA) and 32GB compact flash (3HE06083AA).
- From 7210 SAS 8.0R3 release, SD card is supported with 7210 SAS-K non-ETR platform.
- 7210 SAS-R12 platform is supported from 8.0R3 release.
- 7210 SAS-Mxp non-ETR platform is supported from 8.0R2 release.
- From 8.0R1 release, 7210 SAS-Mxp ETR platform is supported.
- 7210 SAS-R6 supports 8GB compact flash (3HE04708AA) and 32GB compact flash (3HE06083AA) in 7.0R6 release.
- From 7.0R7 release, 7210 SAS-K ETR platform is supported.
- From 7.0R6 release, 7210 SAS-R6 supports following 2nd generation IMM (IMM-SAS-R-b) referred as IMMv2.
 - IMM-SAS-R-b 16TX (P/N: 3HE09156AA).

- From 7.0R4 release, 7210 SAS-R6 supports following 2nd generation IMM (IMM-SAS R-B) referred as IMMv2.
 - IMM-SAS-R-b 10SFP-1SFP+ (P/N:3HE09152AA)
 - IMM-SAS-R-b 2SFP+ (P/N: 3HE09153AA)
 - IMM-SAS-R-b 4SFP+ (P/N: 3HE09154AA)
 - IMM-SAS-R-b 11cSFP (P/N: 3HE09155AA)

NOTE: IMMv2 works only with SF/CPM-b and cannot co-exist with IMMv1 in the same 7210 SAS-R6 chassis

- From 7.0R4 release 7210 SAS-K is supported.
- 7210 SAS-R6 SF/CPM-b (P/N:3HE08154ABRA02) is supported from 7210 SAS 7.0R3 release and it works with 7.0R3 or higher versions of software release.
NOTE: SF/CPM-b 3HE08154ABRA01 and 3HE08154ABRA02 can co-exist in the same 7210 SAS-R6 chassis.
- From release 7.0R1, 7210 SAS-R6 supports only SF/CPM-b (P/N:3HE08154ABRA01). **NOTE:** From 7.0R1 SF/CPM (P/N:3HE08154AARA) is not supported.
- All variants of 7210 SAS-R6, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-D, and 7210 SAS-E.
- Following USB models from Sandisk are supported:
 - Cruzer Fit – size 4GB and 8GB
 - Cruzer Blade – size 4GB and 8GB

3 Changed or Deprecated Commands

3.1 Release 9.0.R1

This section describes the 7210 SAS commands that have been renamed or deprecated in Release 9.0R1.

NOTE: These changes are not applicable to 7210 SAS-E and SAS-D.

3.1.1 Filter Commands

Release 9.0R1 introduces changes to the ACL filter policy commands. The **action** command (with all types and related parameters defining), used to define an action to be performed on a packet matching IPv4/IPv6/MAC ACL policy entry, has been deprecated and replaced by a new **action** command that allows the operator to enter a new CLI context under which individual actions can be configured using **drop**, **forward**, and **action** type commands and their respective parameters.

The operational impact of the **action** command restructuring is as follows:

- Because all command and parameter names were preserved, any ACL configuration prior to Release 9.0R1 remains valid and results in the same configuration result for all cases except the following:

Prior to Release 9.0R1, an operator was able to use CLI or SNMP to configure an IP, IPv6, or MAC filter policy entry's **action** without explicitly specifying the action type. The system would interpret such configuration as "**action drop**". This functionality is no longer supported and action type must always be explicitly specified using new action-type commands under new **action** context. Loading an old configuration file that does not explicitly configure action type will either fail or will result in a different behavior (a filter entry will not be activated as no action is configured). The operator must ensure that any old configuration file always explicitly configures **action** type **drop** for every filter policy entry **action** missing explicit **drop** keyword. During an ISSU upgrade to Release 9.0R1 or newer, the system automatically converts the "**action**" command with no type specified to "**action drop**" for all ACL filter types.

- Starting with Release 9.0R1, **admin save** and **info** commands will save or display filter entry action configuration in a multi-line format (as illustrated below).

- CLI configuration continues to accept a single line format to specify an action with its type and related parameters.

Command prior to Release 9.0R1

```
configure
  filter {ip-filter | ipv6-filter | mac-filter}
    entry
      action
      action drop [optional parameters]
      action forward [optional
        parameters]
```

Command in Release 9.0R1

```
configure
  filter {ip-filter | ipv6-filter | mac-filter}
    entry
      action
        drop
        drop [optional parameters]
        forward [optional parameters]
```

3.1.2 System Commands

7210 SAS cron commands have been aligned to match the changes in 7x50 CLI command structuring for cron commands.

Command prior to Release 9.0R1

```
configure
  cron
```

Command in Release 9.0R1

```
configure
  system
    cron
```

The following PTP command has been renamed to align with the terminology used in the ITU-T profiles: **config>system>ptp>peer>priority** and is changed to: **config>system>ptp>peer>local-priority**

3.1.3 SNMP Commands

The following command has been renamed: **show system security communities** and is changed to: **show system security snmp community [community-name]**

3.1.4 Router Protocols Commands

3.1.4.1 IS-IS Commands

In Release 9.0R1, the **suppress-default** command is changed to **ignore-attached-bit**. The following displays changes in the **configure>router** context

Command prior to Release 9.0R1

```
configure
  system
    isis
      suppress-default
```

Command in Release 9.0R1

```
configure
  system
    isis
      ignore-attached-bit
```

3.1.4.2 BGP Commands

The command for resolution of RFC 3107 BGP label route prefix using tunnels to a BGP next-hop has been deprecated: **config>router>bgp>transport-tunnel {ldp|rsvp-te|mpls}** and is replaced by **config>router>bgp>next-hop-resolution>label-route-transport-tunnel**. When upgrading to Release 9.0, the following values in the configuration file are converted as follows:

Command prior to Release 9.0R1

```
configure
  router
    bgp
      no transport-tunnel
```

Command in Release 9.0R1

```
configure
  router
    bgp
      next-hop-resolution
        label-route-transport-tunnel
          family ipv4
            resolution-filter
            resolution-filter ldp
          family vpn
            resolution-filter
            resolution-filter ldp
```

Command prior to Release 9.0R1

```

configure
    router
        bgp

.....
        transport-tunnel ldp

```

Command in Release 9.0R1

```

configure
    router
        bgp
.....
        family ipv4
            resolution-filter
            resolution-filter
            ldp
        family vpn
            resolution-filter
            resolution-filter
            ldp

```

Command prior to Release 9.0R1

```

configure
    router
        bgp

.....
        transport-tunnel rsvp-te

```

Command in Release 9.0R1

```

configure
    router
        bgp
.....
        family ipv4
            resolution-filter
            resolution-filter rsvp
        family ipv4
            resolution-filter
            resolution-filter rsvp

```

Command prior to Release 9.0R1

```
configure
    router
        bgp
            no transport-tunnel

.....
            transport-tunnel mpls
```

Command in Release 9.0R1

```
configure
    router
        bgp
            .....
                family ipv4
                    resolution-filter
                    resolution-filter ldp
                    rsvp
                family ipv4
                    resolution-filter
                    resolution-filter ldp
                    rsvp
```

3.1.5 MPLS Commands

3.1.5.1 MPLS Commands

The following command has been deprecated: **config>router>mpls-labels>static-labels max-lsp-labels max-lsp-labels max-svclabels max-svc-labels** and is replaced by: **config>router>mpls-labels>static-label-range static-range**. The static LSP and service label ranges are now collapsed into a single range usable by all applications requiring a static label. During the upgrade to Release 9.0, the configuration file is updated to reflect the new single **static-label-range** CLI with a range size matching the value saved previously under the **max-svc-labels** parameter.

3.1.5.2 LDP Commands

In Release 9.0R1, the LDP CLI has new commands, updates to existing commands, and deprecations of commands. Changes to the CLI hierarchy also impact the configuration of LDP IPv4 peers and targeted sessions. The following is a summary of the major changes.

Command prior to Release 9.0R1

```

configure
  router
    ldp
      peer-parameters
        peer-ip-address

```

Command in Release 9.0R1

```

configure
  router
    ldp
      session-parameters
        peer-ip-address

```

Command prior to Release 9.0R1

```

configure
  router
    ldp
      peer-parameters
        peer-ip-address

```

Command in Release 9.0R1

```

configure
  router
    ldp
      tcp-session-parameters
        peer-transport ip-address

```

Command prior to Release 9.0R1

```
configure
  router
    ldp
      interface-parameters
```

Command in Release 9.0R1

```
configure
  router
    ldp
      interface-parameters
        ipv4
```

Command prior to Release 9.0R1

```
configure
  router
    ldp
      interface-parameters
        interface ip-int-name
```

Command in Release 9.0R1

```
configure
  router
    ldp
      interface-parameters
        interface ip-int-name [dual-
          stack]
          ipv4
```


Command prior to Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name
```

Command in Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name [dual-
                    stack]
```

Command prior to Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name
                    bfd-enable
```

Command in Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name
                    bfd-enable
```

Command prior to Release 9.0R1

```
configure
    router
        ldp
            targeted-session
```

Command in Release 9.0R1

```
configure
    router
        ldp
            targeted-session
                ipv4
```

Command prior to Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name
                    multicast-traffic
                        {enable|dis-able}
```

Command in Release 9.0R1

```
configure
    router
        ldp
            interface-parameters
                interface ip-int-name
                    ipv4
                        fec-type-capability
                            p2mp-ipv4
                                {enable|disa
                                    ble}
```

The following details the changes to the syntax of the show commands to better convey the scope of applicability of the command.

Command prior to Release 9.0R1

```
show
    router
        ldp
            peer
```

Command in Release 9.0R1

```
show
    router
        ldp
            targ-peer
```

Command prior to Release 9.0R1

```
show
    router
        ldp
            peer-parameters
```

Command in Release 9.0R1

```
show
    router
        ldp
            session-parameter
```

Command prior to Release 9.0R1

```
show
    router
        ldp
            peer-template
```

Command in Release 9.0R1

```
show
    router
        ldp
            targ-peer-template
```

Command prior to Release 9.0R1

```
show
    router
        ldp
            peer-template-map
```

Command in Release 9.0R1

```
show
    router
        ldp
            targ-peer-template-map
```

Command prior to Release 9.0R1

```
show
    router
        ldp
            auth-keychain
```

Command in Release 9.0R1

```
show
    router
        ldp
            tcp-session-parameters [keychain
                                   keychain]
```

Command prior to Release 9.0R1

```
show
  router
    ldp
      bindings
        prefix ip-prefix|ip-prefix-
          length
```

Command in Release 9.0R1

```
show
  router
    ldp
      bindings
        prefixes prefix ip-prefix|ip-
          prefix-length
```

Command prior to Release 9.0R1

```
show
  router
    ldp
      bindings
        active
```

Command in Release 9.0R1

```
show
  router
    ldp
      bindings
        active
          prefixes prefix ip-
            prefix|ip-prefix-
              length
```

Command prior to Release 9.0R1

```

show
  router
    ldp
      status

```

Command in Release 9.0R1

```

show
  router
    ldp
      statistics

```

Command prior to Release 9.0R1

```

show
  router
    ldp
      bindings

```

Command in Release 9.0R1

```

show
  router
    ldp
      bindings

```

The following is the numbered explanation:

1. Command shows parameters configured under **config>router>ldp>targ-session>peer** *x.y.z.w*.
2. Command shows parameters configured under **config>router>ldp>session-parameters>peer** *x.y.z.w*.
3. Command shows parameters configured under **config>router>ldp>targ-session** **peertemplate**.
4. Command shows parameters configured under **config>router>ldp>targ-session>peer-template-map**.
5. The output of the **show router ldp auth-keychain** command can now be displayed using **show router ldp tcp-session-parameters** [**keychain** *ke chain*]
6. The **show router ldp bindings prefix** command has been moved one level deeper in the CLI hierarchy in a new **prefixes** CLI context.
7. A new **prefixes** level in the CLI hierarchy has been added to the **show router ldp binding** context to display the unicast IPv4 FECs.
8. Some commands related to statistics, previously displayed under the **show router ldp status** context, have now been moved to the new **show router ldp statistics** context.

The following changes to the syntax of the clear and debug commands were made to better convey the scope of the applicability of the command.

Command prior to Release 9.0R1

```
clear
      router
        ldp
          instance
```

Command in Release 9.0R1

```
clear
      router
        ldp
          instance family [ipv4]
```

Command prior to Release 9.0R1

```
clear
      router
        ldp
          interface ip-int-name
```

Command in Release 9.0R1

```
clear
      router
        ldp
          interface ip-int-name [family
                                [ipv4]
```

Command prior to Release 9.0R1

```
debug
    router
        ldp
            interface interface-name
```

Command in Release 9.0R1

```
debug
    router
        ldp
            interface-name [family [ipv4]]
```

3.1.6 Layer-2 Services Commands

3.1.6.1 VPLS Services Commands

The following command has been deprecated: **config>service>vpls>allow-ip-int-binding** and is replaced by: **config>service>vpls>allow-ip-int-bind**.

The latter is a CLI node, allowing the future addition of sub-commands below it.

3.1.7 Layer-3 Services Commands

The command for resolution of a VPN-IPv4 or VPN-IPv6 prefix to a BGP next-hop has been renamed: **config>service>vprn>auto-bind {ldp|rsvp-te|mpls|}** and is changed to **config>service>vprn>auto-bind-tunnel**. When upgrading to Release 9.0, the following values in VPRN auto-bind in the configuration file will be converted as follows:

Command prior to Release 9.0R1

```
configure
    service
        vprn service-id auto-bind
```

Command in Release 9.0R1

```
configure
    service
        vprn service-id
            auto-bind-tunnel
                resolution filter
```

Command prior to Release 9.0R1

```
configure
    service
        vprn
            auto-bind ldp
```

Command in Release 9.0R1

```
configure
    service
        vprn
            auto-bind-tunnel
                resolution filter resolution-
                    filter ldp
```

Command prior to Release 9.0R1

```
configure
    service
        vprn
            auto-bind mpls
```

Command in Release 9.0R1

```
configure
    service
        vprn
            auto-bind-tunnel
                resolution filter resolution-
                    filter ldp rsvp
```

Command prior to Release 9.0R1

```
configure
    service
        vprn
            auto-bind mpls-gre
```

Command in Release 9.0R1

```
configure
    service
        vprn
            auto-bind-tunnel
                resolution filter resolution-
                    filter gre ldp rsvp
```

Command prior to Release 9.0R1

```
configure
  service
    vprn
      auto-bind rsvp-te
```

Command in Release 9.0R1

```
configure
  service
    vprn
      auto-bind-tunnel
        resolution filter resolution-
          filter rsvp-te
```

3.1.8 OAM Commands

3.1.8.1 Diagnostic Commands

As of Release 9.0R1, legacy and Alcatel-Lucent-specific OAM commands no longer support the **send-control** option. Operators should stop using the **send-control** option with the following OAM functions for both interactive CLI, SNMP and SAA:

- **cpe-ping**
- **mac-ping**
- **mac-populate**
- **mac-purge**
- **mac-trace**

With Release 9.0R1, the **send-control** option will no longer be available. All SAA tests that include this option will fail to start as of Release 9.0R1. This option must be removed from the SAA tests.

4 Software Upgrade Notes

4.1 Upgrade to 9.0R1 or later releases

4.1.1 Software Upgrade to 9.0R1 or Later Release from previous Releases



Note: 7210 SAS-K means 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C platforms. 7210 SAS-Sx refers 7210 SAS-Sx and 7210 SAS-S platforms unless specified.

Previous 7210 releases boot.tim is not compatible with 9.0R1 or later both.tim (of SAS-M, SAS-X, SAS-T, SAS-K, SAS-Mxp, SAS-Sx) and cpm.tim, iom.tim (of SAS-R). There is no compatibility issue in case of SAS-E and SAS-D. Two procedures are explained for software upgrade

4.1.1.1 Software upgrade using auto init

Before Upgrade to 9.0R1 or later software from older software versions using auto boot, copy 9.0R1 or later boot.tim in flash of node and both.tim in server. Delete bof.cfg in the node and upgrade to new software using CLI “admin reboot upgrade”.

4.1.1.2 Manual software upgrade

Before Upgrade to 9.0R1 or later software from older software versions, copy 9.0R1 or later boot.tim and both.tim, cpm.tim, iom.tim in flash of node. Then execute CLI “admin reboot upgrade”.

4.1.2 7210 SAS-R “global-res-prof”

Before 9.0R1 release of 7210 SAS-R in “*config>system>global-res-prof>sf-ingress-internal-tcam*” policy for “g8032-fast-flood” and “eth-cfm-upmeps” odd or even values were allowed for configuration. From 9.0R1, only odd values (1, 3, 5, or 7) can be configured for “g8032-fast-flood” and only even values (0, 2, 4 or 6) can be configured for “eth-cfm-upmeps”.

During upgrade to 9.0R1 or later version, software will adjust these values accordingly from configuration file. In rare cases if configuration execution fails during upgrade, manually configuration should be changed.

4.1.3 Routing Policies

From Release 9.0 R1 onwards, the use of a community, **as-path**, **as-path-group** or **prefix-list** name starting and ending with '@' is no longer allowed. @...@ is used as identification for parameters being used in policies. Upgrading from a pre-Release 9.0R1 to Release 9.0R1 or higher will mean that configuration files containing such names will fail to execute.

4.2 Upgrade to 8.0R1 or later releases

- In 8.0R1 release following accounting numbers are modified, new numbers are shown in () netInIngressOct(101), netInIngressPkt(102), combinedNetInIngress(103), accessEgressPkt(104), accessEgressOct(105), combinedAccessEgress(106), combinedNetEgress(107), combinedSvcEgress(108), combinedSvcInEgPkt(109), combinedNetInEgPkt(110)
- In case of 7210 SAS-D ETR, 7210 SAS-K, 7210 SAS-M (access-uplink and network mode), 7210 SAS-X, after upgrade to 8.0R1 and if PTP is configured on the system Y.1731 OAM tools 2-DM, 1-DM and SLM start using PTP time stamps. Same behavior exists for 7210 SAS-T (access-uplink and network mode), 7210 SAS-R6 with previous releases.

4.3 Upgrade to 7.0R4 or later releases

- For detailed procedure for upgrading IMMv1 to IMMv2 and bringing up system with IMMv2 7.0R4 or later version of software, refer 7210 SAS Installation Guide.
- From 7.0R4, "*max-ipv6-routes*" CLI is available under "*config>system>res-profile>router*".

4.4 Upgrade to 7.0R1 or later releases

4.4.1 UPGRADE TO SF/CPM-B (P/N:3HE08154ABRA)

Platforms applicable: 7210 SAS-R6

From release 7.0R1, 7210 SAS-R6 supports only SF/CPM-b (P/N:3HE08154ABRA).

7210 SAS-R6 SF/CPM (P/N:3HE08154AARA) and SF/CPM-b (P/N:3HE08154ABRA) cannot co-exist in a single chassis. To upgrade to SF/CPM-b execute the following:

Step 1. Copying required version of the software image using the below methods.
Copy 7.0R1 or greater versions of software and config files to cf2: of existing SF/CFM, modify bof.cfg to point to this software. Take out flash cf2: from SF/CFM and put in to SF/CFM-b.

OR

Copy the required images (7.0R1 or higher version) and config files to flash of SF/CPM-b, make sure to bof.cfg point to the required software and config file.

Step 2. Power down the chassis.

Step 3. Remove both (if in use) the SF/CPM and replace both with SF/CPM-b.

Step 4. Plug-in the appropriate connectors (for example: console port connection) to the new SF/CPM-b cards.

Step 5. Power up the chassis.

Step 6. On boot up with SF/CFM-b, "show card" output will display "cpm-sf-b-sas-R6" the correct names for the SF/CPM-b for slot A and B.

NOTES:

- No configuration changes are needed to use SF/CPM-b.
- The upgrade is service affecting.

4.5 Upgrade to 6.0R6 or later releases

4.5.1 FILTERS

Platforms applicable: 7210 SAS-E

In release 6.0R6, the ACL TCAM allocation scheme has been modified. As part of this change, software allocates only one (1) additional entry for every group of resources allocated for use of ACLs. In previous releases configuration if all ACL entries are used up, it is required to remove one of IP, MAC, or IPv6 entry before upgrading to 6.0R6 release.

4.6 Upgrade from 6.0R4 or prior releases

4.6.1 1830 VWM (CWDM) MANAGEMENT

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E

The 7210 SAS release 6.0R5 adds supports for provisioning of cards inserted into the slots available on the 1830 VWM devices. The user must provision the card and card-type (also known as, module type) before the card can be managed by the 7210 SAS after upgrade to 6.0R5 release.

4.7 Upgrade from 3.0 releases

4.7.1 NETWORK QOS POLICY

Platforms applicable: 7210 SAS-M (Network mode) and 7210 SAS-X

During upgrade from 3.0 to 4.0 or 5.0 releases following was true:

- For each "network" qos policy of "ip-interface" type in the configuration file, system generated "mpls-lsp-exp-profile-map" policy and was attached to network qos policy. mpls-lsp-exp-profile-map policy id was equal to network qos policy id. "mpls-lsp-exp-profile-map" policy was populated with "lsp-exp <lsp-exp-value> profile {in|out}" information. lsp-exp and profile values were copied from corresponding network qos policy of the configuration file.
In case of upgrade from 3.0 to 6.0 or later releases
- Following console message appears "MINOR: CLI MPLS LSP EXP profile should be configured through the profile map."
- "profile in|out" configured in "network" qos policy of 3.0 configuration file will be ignored during upgrade. This means user defined 3.0 "profile" definition is lost after upgrade to 6.0 or later releases.
- System attaches default, that is, "mpls-lsp-exp-profile-map" policy id 1 to network qos policy.

Workaround:

- First upgrade from 3.0 to 4.0 or 5.0, save config file then upgrade to 6.0 or later releases.

OR

- After upgrade to 6.0 or later release, user need to configure "mpls-lsp-exp-profile-map" policy and attach to "network" qos policy.

4.8 Upgrade from 4.0 or 5.0 releases

4.8.1 CLI

Platforms applicable: 7210 SAS-M (Network mode) and 7210 SAS-X

- After upgrade to release 6.0 or later releases, system defaults to "ldp-local-fc-enable", CLI "config>qos>ldp-local-fc-enable" is not available from release 6.0.
- Profile parameter [profile {in|out}] under all the Eth-CFM SAA tests (Loopback, Linktrace, 2DM, and 2SLM) was not supported in previous releases and CLI commands has been removed in release 6.0R1. During upgrade profile information, if existed in the config file, will be ignored.

4.9 Upgrade from 5.0R1 or 5.0R2 releases

4.9.1 LAG CONFIGURATION

Platform applicable: 7210 SAS-M and 7210 SAS-X

The 7210 node is allotted a fixed amount of MAC addresses during manufacturing. The base address and the number of MAC addresses is specified on the back of the chassis on the chassis label and also shown in the show chassis command. On 7210 SAS-M or SAS-X, software reserves about 28 addresses for its use to assign MAC addresses to all the ports, system mac etc. MAC addresses are assigned to the LAG using the LAG ID as the offset.

With this allocation scheme, if the total number of MAC addresses is, for e.g 44, then from LAG ID 17 up to the maximum amount of LAG configured, user needs to assign MAC address statically before upgrade to 5.0R3 release, otherwise upgrade to 5.0R3 or later release may fail if system does not have required number of MAC addresses.



Note: The issue will be seen only on nodes which do not have enough MAC addresses, that is, those less than 53 addresses.

From 5.0R3 or later release "show chassis" CLI displays Base MAC address and number of MAC's assigned to node.

4.9.2 DOT1X TUNNELING

Platforms applicable: 7210 SAS-M Network mode and 7210 SAS-X

When "dot1x tunneling" configured on access port, with release 5.0R1 or R2 "admin save" saved "dot1x tunneling" in wrong context of config file. This resulted in errors during next re-boot of node.

Workaround is to modify the saved config so that "dot1x tunneling" appears after "mtu".

Example shown below:

```
port 1/1/1
  ethernet
    mode access
    access
    exit
    mtu 9212
    dot1x
    tunneling
  exit
```

4.10 Upgrade to 5.0R1 or later releases from prior releases

4.10.1 ACCOUNTING RECORD NUMBERS

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

In 5.0R1 release following accounting numbers are modified, new numbers are shown in () netInIngressOct(52), netInIngressPkt(53), combinedNetInIngress(54), accessEgressPkt(55), accessEgressOct(56), combinedAccessEgress(57), combinedNetEgress(58), combinedSvcEgress(59), combinedSvcInEgPkt(60), combinedNetInEgPkt(61)

4.10.2 NON SUPPORTED CLI

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

In previous releases 7210 SAS was allowing to configure non supported feature CLI's. In 5.0R1 some of the non supported CLI's are removed. In previous releases if user had configured non supported feature CLI, configuration will error out during upgrade to 5.0R1 release. It is recommended to check for non supported CLI by loading config file with 5.0R1 and remove non supported CLI's from config file before final upgrade.

4.10.3 TACACS+

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

From 5.0R1 tacacs+ "single-connection" option is deprecated. During upgrade following warning message displayed



Note: CLI Line:xx "single-connection" This command has been deprecated.

4.10.4 SPLIT HORIZON GROUP NAME

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

During upgrade to 5.0R1 or later release, if "split-horizon-group <group-name>" CLI, where <group-name> configured with name having spaces, config file execution will error out during upgrade.

Workaround is to edit the config file before upgrading to ensure split horizon group names are with double quotes or no space(s) in the group name.

4.10.5 LOGGING

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

In 5.0 or later releases , as part of log-id config (configure>log>log-id <num>) needs to have "from change".

This is necessary for the box to generate config change logger/trap messages, modify configuration file accordingly.

4.11 Upgrade to 4.0R4 or later release from prior releases

4.11.1 SAP INGRESS QOS

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

During upgrade to 4.0R4 or later release, the following warning message displayed can be ignored:

MAJOR: CLI #1010 Saps in the system will be re-configured without Sap Indexs, because SAP index file could not be located.

After upgrade to 4.0R4 or higher, it is recommended to save the configuration. When the configuration is saved with 4.0R4 or higher build, sap index file with extension .sdx is automatically generated and saved at a location where the configuration file is stored.

4.11.2 Y.1731 MA NAME

Platforms applicable: 7210 SAS-M, 7210 SAS-X, 7210 SAS-E, 7210 SAS-D

From 4.0R1, Y1731 MA name should be unique across the system. During upgrade to 4.0R1 or Later releases if system finds duplicate names in the configuration file then upgrade to new configuration fails. Before upgrade, it is recommended to modify the configuration file.

Example CLI:

"association 1 format icc-based name "abcdabcdabcd1", name "abcdabcdabcd1" should be unique across system for successful upgrade to 4.0R1 or higher.

4.12 Upgrade to 1.1R7 or later releases from prior releases

4.12.1 SERVICES

Platforms applicable: 7210 SAS-M

In 1.1R7, a default SAP and a dot1q SAP cannot be configured along with enabled egress filters when the SAPs are configured on the same port.

If such a configuration exists in the startup-config, it is recommended that the following procedure is used before performing an upgrade:

- Step 1.** Create a config file on the cf1 flash which contains the CLI commands to provide basic in-band connectivity and management functions.
- Step 2.** Add the following command in the startup configuration file. Note that the config-file parameter is the file mentioned in step #1. `config>system>bootbad-exec "cf1:./<config-file>"`
- Step 3.** Follow the [Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, Mxp, and Sx \(standalone mode\) devices](#).

4.13 Upgrade to 1.1R6 or later releases from prior releases

4.13.1 ACL

Platforms applicable: 7210 SAS-M

In 1.1R6, number of egress ACLs is restricted to 256 for combined IP and MAC criteria. If the 7210 SAS M is booted with a configuration file containing more than 256 egress ACLs, the configuration will error out. It is recommended that, in such cases, the following procedure is used before performing an upgrade.

- Step 1.** Create a config file on the cf1 flash which contains the CLI commands to provide basic in-band connectivity and management functions.

- Step 2.** Add the following command in the startup configuration file. Note that the config-file parameter is the file mentioned in step #1. config>system>bootbad-exec "cf1:./<config-file>"
- Step 3.** Follow the [Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, Mxp, and Sx \(standalone mode\) devices](#).

5 Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, Mxp, and Sx (standalone mode) devices

**Note:**

Refer the SR 7x50 documentation for 7210 SAS-Sx (satellite mode) software upgrade procedure and known issues.

The following sections contain information for upgrading to the 9.04 software version. In particular, there are sections that describe the following:

- Standard Software Upgrade Procedure
- Procedure for performing a standard, service-affecting upgrade.

5.1 Standard Software Upgrade Procedure

This section describes the standard software upgrade procedure which is service-affecting:

- Each software release includes a boot loader (boot.tim) and the software image (both.tim).
- The boot loader initiates the loading of the 7210 SAS OS image based on the boot options file (bof.cfg) settings.

The following steps describe the software upgrade process:

Step 1. Backup 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:**

- Alcatel-Lucent recommends making backup copies of the BOOT loader (boot.tim), software image (both.tim) and configuration files, should reverting to the old version of the software be required.

Step 2. Copy 7210 SAS images to cf1:

The 7210 SAS image files both.tim must be copied to the cf1: device on the 7210 SAS node. It is good practice to place all the image files for a given release in an appropriately named subdirectory off the root, for example, cf1:\9.0R4 Copying the boot.tim and other files in a given release to a separate subdirectory ensures that all files for the release are available should downgrading the software version be necessary.

Note: Applicable only to 7210 SAS-D Devices



- The 7210 SAS-D of 64MB flash (part numbers 3HE05676AAAA01, 3HE05676ABAA01, 3HE05677AAAA01, 3HE05677ABAA01 and 3HE06537AAAA01) can accommodate one set of boot.tim and both.tim files, the users are required to overwrite existing files with new files in cf1. With the Enhanced 7210 SAS-D (SAS-D with 128MB flash) supported from 4.0R6, more than one set of boot.tim and both.tim files can be stored on flash.

Step 3. Copy boot.tim to the root directory on cf1:

The BOOT Loader file is named boot.tim. This file must be copied to the root directory of the cf1: device.

Note:



- If it is not possible to overwrite cf1:\boot.tim file, Change the cf1:\boot.tim attributes using **file attrib -r cf1:\boot.tim** command.
- In case of 7210 SAS-T, 7210 SAS-K, 7210 SAS-Mxp, 7210 SAS-Sx and other supported drives such as "cf2:" or "uf1:" can also be used for storing boot.tim image. Note that the valid "bof.cfg" file is in the same drive from where the boot.tim is used. Order of search for boot loader is cf1:/boot.tim, cf2:/boot.tim, uf1:/boot.tim. bof.cfg is read from the drive where boot.tim is loaded.

Step 4. Modify the boot options file to boot the new Image.

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

Step 5. When upgrading to 9.0R4, execute the **admin reboot upgrade** command. Note that, executing '**admin reboot upgrade**' command system will upgrade the bootrom if required.

Allow the boot sequence to complete and verify that the card comes up.

Step 6. Applicable only to 7210 SAS-D, E, M, and X devices

Upgrade the Golden BOOT Loader (only if all of the above steps were successful).

After successfully booting of the new version of 7210 SAS image, upgrade the golden boot loader by executing the **admin update-golden-bootstrap cf1:/boot.tim** command.

Note:

After upgrade to 4.0 or later software images, during next boot, if the user forgets the BOF password and fails to provide a correct password, after three attempts, the system prompts the user to reset the BOF password to factory default. If user accepts BOF password recovery, as a security measure, the system also resets the flash to factory defaults (that is, it removes all the files from the flash except the boot image file (cf1:\boot.tim) and Timos image file (cf1:\both.tim)) and reboots the node with the factory default settings.

Refer the 7210 SAS User Guides and "**BOF PASSWORD RECOVERY**" in [Enhancements](#) section of this document for more details.

6 Software Upgrade Procedures for 7210 SAS-R

Note: 7210 SAS-R means 7210 SAS-R6 and 7210 SAS-R12 platforms.

The following sections contain information for upgrading to the 9.0R4 software version. In particular, there are sections that describe the following:

- Standard Software Upgrade Procedure
Procedure for performing a standard, service-affecting upgrade.

6.1 Standard Software Upgrade Procedure

This section describes the standard software upgrade procedure which is service-affecting:

- Each software release includes a boot loader (boot.tim) and the software image (cpm.tim and iom.tim).
- The boot loader initiates the loading of the 7210 SAS OS image based on the boot options file (bof.cfg) settings.

The following steps describe the software upgrade process:

Step 1. Backup 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:

- Alcatel-Lucent recommends making backup copies of the BOOT loader (boot.tim), software image and configuration files, should reverting to the old version of the software be required.

Step 2. Copy the 7210 SAS-R OS Images to cf2:

Note: Software images are common for both 7210 SAS-R6 and 7210 SAS-R12 platforms.

The 7210 SAS-R image files must be copied to the cf2: device on the 7210 SAS-R node. It is a good practice to place all the image files for a given release in an appropriately named subdirectory off the root, for example, cf2:\9.0R4.

Copying the boot.tim and other files in a given release to a separate subdirectory ensures that all files for the release are available for downgrading the software version to be necessary.

Note: The 7210 SAS-R drives such as "cf1:" or "uf1:" can also be used for storing boot.tim image. Note that the valid "bof.cfg" file should be in the same drive from where the boot.tim is used. Order of search for boot loader is cf1:/boot.tim, cf2:/boot.tim, and uf1:/boot.tim. bof.cfg is read from the drive where boot.tim is loaded.



Note:

- If it is not possible to overwrite the cf2:\boot.tim file, change the cf2:\boot.tim attributes using file **attrib -r cf2:\boot.tim** command.

Step 3. Copy boot.tim to the Root Directory on cf2:

The BOOT Loader file is named boot.tim. This file must be copied to the root directory of the cf2: device.

Step 4. Modify the Boot Options File to Boot the New Image

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

Step 5. [Redundant CPMs or CFMs] Synchronize Boot Environment

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

Step 6. Reboot the Chassis

When upgrading to 7210 SAS 9.0R4, execute the admin reboot upgrade command. Note that, executing '**admin reboot upgrade**' command in the system upgrades the bootrom.

Step 7. Verify the Software Upgrade

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

Note: If any card fails to occur online after the upgrade, contact the Nokia Technical Assistance Center for information on corrective actions.

It is recommended to save the configuration “**admin save**” after an upgrade has been performed and the system is operating as expected. This ensures that all configuration is saved in a format that is fully compatible with the newly running release.

Note:

During next boot, if the user forgets the BOF password and fails to provide a correct password, after three attempts, the system prompts the user to reset the BOF password to factory default. If user accepts BOF password recovery, as a security measure, the system also resets the flash to factory defaults (that is, it removes all the files from the flash except the boot image file (cf2:\boot.tim) and Timos image file) and reboots the node with the factory default settings.

7 Resolved Issues

**Note:**

- 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

7.1 Resolved in R9.0R4

The following are specific technical issues that have been resolved in Release 9.0R4 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

- In case of 7210 SAS-R, 7210 SAS-Mxp and 7210 SAS-T, due to a counter roll-over, ingress meter statistics were not shown properly when it crosses a limit. This is resolved in 7210 SAS 9.0R4 release. [244799,249244-MI]
- If an in-band SSH session is idle for around 1 minute, the first CLI command executed can be unresponsive for a few seconds. [250103-MI]
- Port loopback with mac-swap was failing for frame size more than 1500 if the port encapsulation is either dot1q or qinq. This issue is resolved in 7210 SAS 9.0R4 release. [250218-MI]
- When remote mirror is configured, global revert of LSPs to a new path generates few internal trace messages. This issue is resolved in 7210 SAS 9.0R4 release. [248096-MI]
- In case of 7210 SAS-K 2F4T6C, CCM not was converged on standby spoke sdp MEP configured on VPLS service. This issue is resolved in 7210 SAS 9.0R4 release. [247446-MI]
- In case of 7210 SAS-K, PTP was not proper if SFP port bounces where PTP packets are received. Workaround was to shut no shut of that particular port and issue was not seen with copper port. The 7210 SAS 9.0R2 resolves issue with SFP port. [241704-MI]
- 7210 SAS-K 2F4T6C supports ETH-CFM, G8032, MPLS OAM tools, Y.1564, OAM-PM, Twamp-lite, Twamp server from 7210 SAS 9.0R4 release.
- 7210 SAS-Sx (standalone mode) supports BFD from 7210 SAS 9.0R4 release.

7.2 Resolved in R9.0R3

The following are specific technical issues that have been resolved in Release 9.0R3 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

- In previous release, using an unsupported svcType value “2” while creating a service from SNMP used to reset the system, this issue is resolved. [244584-MA]
- If 7210 SAS-R6 receives a constant burst of SNMP request packets to IF-MIB table, sometimes few request packets were dropped resulting in snmp time out or slow response from node. This issue is resolved in this release. [235593-MI]
- In previous release of 7210 SAS-Mxp ping RTT values were high if PTP is configured and PTP is not yet locked. This issue is resolved now. [240288-MI]
- With previous release some times system reboot was observed when eth-cfm MEP were configured on spoke-sdp in scaled number. This is resolved now. [240686-MA]
- With previous release of 7210 SASR some times it was observed that traffic fails to egress out of the Eth-ring sap, shut/no shut of Eth-ring used to restore traffic. This issue is resolved now. [246057-MI]
- In case of 7210 SAS-M, SAS-X, SAS-D and SAS-K, an unexpected system reset could have occurred when IEEE 1588 PTP was operating over an extended period of time. This reset would have occurred after 776 days of continuous IEEE 1588 PTP operation with default message rates of 64 packets per second, or after 388 days if any external IEEE 1588 peers or ports were using message rates of 128 packets per second. In prior releases, the system reset can be prevented by taking following action prior to the expiration of above mentioned periods:
 - Remove all PTP peer and port configuration, shut down PTP with configure system ptp shutdown, remove all dynamic peer information with clear system ptp inactive peers and then reconfigure and re-enable PTP. This issue has been resolved. [229549-MA]

7.3 Resolved in R9.0R2

The following are specific technical issues that have been resolved in Release 9.0R2 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

- For 7210 SAS-R, IPv6 neighbor configuration disappeared after CPM switch over from active to backup and again to active in previous release, this issue is resolved. [240528-MI]
- Deletion of primary address on routing interface when secondary address is configured removed all addresses, this issue is resolved. [240071-MI]
- With previous release of 7210 SAS-R Traffic drops were seen when MBB takes place from the FRR path to the primary path in an inter IMM case (one-to-one FRR being used). This issue is resolved. [237416-MA]
- In case of 7210 SAS-K with previous release, PTP was not proper if SFP port bounces where PTP packets are received. Workaround was to shut no shut of that particular port. This issue is not seen with copper port. This issue is resolved. [241704-MI]
- In previous releases at a very high traffic rate few packet drop observed during arp-refresh on the services running over the router interface, this issue is resolved now. [238090-MI]
- In previous releases, DSCP value in the SDP keep-alive reply packet sent by 7210 was set to 0x00 instead of 0xE0, this anomaly is corrected. [239547-MI]
- With previous release, In case of 7210 SAS-Sx, VPRN service FRR switchover times were greater than 50ms, this issue is resolved. [226453-MA]
- In case of 7210 SAS-R and 7210 SAS-T – In a MPLS-TP scenario, lsp ping/trace was failing for the associated channel having a value of “none” when un-numbered mpls-tp interfaces were used, this issue is resolved. [241611-MI]
- In case of 7210 SAS-R, Epipe VLAN-range LAG SAP statistics were not incrementing if shut/no shut is done on the primary port of the lag, this issue is resolved now. [231797-MI]
- In case of 7210 SAS-K with previous release with scaled configuration EFM OAM sessions were flapped, this issue is resolved. [241307-MI]

7.4 Resolved in R9.0R1

The following are specific technical issues that have been resolved in Release 9.0R1 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

- For 7210 SAS-R a rare occurrence of IMM is not booting up post chassis power cycle is observed in previous releases. This issue is resolved. [229448-MA]

-
- In case of 7210 SAS-K If at least one SAP belong to same port is classified using ipv6-criteria in sap-ingress QoS policy and other SAPs are classified using mac-criteria in the same or different sap-ingress QoS policy, packets that are supposed to follow mac-criteria and classified to an FC, can be wrongly classified to a different FC. This issue is resolved. [237292-MI]
 - In case of 7210 SAS-R, Multiple SAPs on a LAG are created in a particular order, R-VPLS SAP with lower vlan tag value, and then an IES sap with vlan tag value higher than R-VPLS sap. With this config, now adding a new LAG member port that belongs to a different IMM than the current member, then the system reboots in previous releases. This issue is resolved now. [238721-MA]
 - In previous releases, configuring an unsupported command, “*interface*” under VPLS service is causing system reboot. This issue is resolved in this release. [239520-MI]
 - In previous release, 7210 SAS-Sx (standalone mode) Eth-Ring switch over times were greater than 50ms when the number of data services exceed 40 (approx.120 SAPs), this has been improved in 9.0R1 release. [227426-MI]
 - In case of 7210 SAS-R, 7210 SAS-Mxp, 7210 SAS-T, and 7210 SAS-Sx, CFM tests (LB, LT, DMM, SLM through SAA or On-Demand) between UP MEPs were failing when UP MEP is configured on VPLS/EPIPE SAP, Spoke-SDP uses LDP as transport LSP and ECMP is configured on LER with multiple paths to remote PE. This issue is resolved. [223231-MI]
 - In case of 7210 SAS-R and 7210 SAS-T, If a bfd-template was created without ‘*begin*’ and ‘*commit*’ commands in previous releases and saved, it was saved incorrectly. If this configuration was used as bootup configuration, configuration errors out during system bootup. From 9.0R1 configuration is properly saved.

8 New Features

8.1 Release 9.0R4

The following items describe features added to 9.0R4 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

8.1.1 7210 SAS Platform – 7210 SAS-S 1/10GE fiber variants in standalone mode

This release includes support for the following 7210 SAS-S 1/10GE variants fiber variants - 7210 SAS-S 48F 4SFP+ AC, 7210 SAS-S 48F 4SFP+ DC, 7210 SAS-S 24F 4SFP+ AC and 7210 SAS-S 24F 4SFP+ DC. These are 1RU platform, providing an option of using either 48 x 100/1000Mbps SFP ports or 24 x 100/1000Mbps SFP ports. In addition, it has 4 x 10GE SFP+ ports for use as uplinks.

Both these variants provide the following support (unless specified otherwise):

- 1RU platform.
- Front access to all the 1GE and 10GE service interfaces.
- Line-rate switching on all the ports.
- One integrated power supply AC or DC -48V and one hot-swappable power supply (AC, DC -48V).
- Non-removable Fans on the rear of the chassis
- Console port and Out-of-band Ethernet management interface on the rear front panel of the chassis

This release supports only standalone mode of operation. In standalone mode of operation, 7210 SAS-S 1/10GE platform is a full featured IP/MPLS router suitable for use in 1GE access and access aggregation network to aggregate up to 48 x 1GE ports/customers onto 4 x 10GE uplinks. In standalone mode, the device needs to be managed independently by assigning an IP address. It can be managed using the console, or using any of the 1GE or 10GE ports (in-band management) or using the out-of-band Ethernet management port. It supports CLI command interface and SNMP for device management. All the service and service attributes configuration (For example: QoS policies, ACLs, accounting logs, etc.), along with control plane protocol configuration (For example: RSVP, IS-IS, etc.) is done on the 7210-

standalone device. It provides operators to use IP/MPLS-based transport mechanisms for providing highly available and resilient services. It can support both Layer-2 VPN services and Layer-3 VPN services with per SAP hierarchical ingress policing and per port hierarchical egress shaping/scheduling, along with extensive Ethernet and MPLS OAM support. In release 7210 SAS 9.0R4, the service and protocol feature support on SAS-S 1/10GE platform is compatible with 7210 SAS-Sx 1/10GE platform.

For more information refer to the all the 7210 SAS-M, Mxp, Sx, T, X, R6, R12 User Guides.

8.1.2 7210 SAS Platform – 7210 SAS-S 10/100GE 64SFP+ 4CFP4 in satellite mode – BETA only.

This release includes supports the 7210 SAS-Sx 10/100GE 64SFP+ 4CFP4, which can be used for high density 10GE port aggregation to 100GE uplinks. It supports 64 x 1GE/10GE SFP+ ports and 4 100GE CFP4 ports. This release supports only satellite mode of operation.

NOTE: Satellite mode support is BETA and cannot be used in production networks.

Some of major features of the platform are:

- 1.5RU platform designed for NEBS compliance.
- Front access to all the 10GE and 100GE service interfaces.
- Line-rate switching on all the ports for IMIX traffic.
- Redundant hot-swappable power supply (AC, DC -48V).
- Non-removable Fans on the rear of the chassis.
- Console port and Out-of-band Ethernet management interface on the rear of the chassis.

NOTES:

- User has an option to configure the role of the chassis using the BOF menu to configure the BOF parameter 'chassis-role' and 'host-type'. The BOF parameter 'chassis-role' allows the user to nail down the chassis to operate in either satellite mode or standalone mode of operation (In release 7210 SAS 9.0R4 only satellite mode support is available. Standalone mode can be used when support is available in future release). For more details on use of this parameter use the 7210 SAS M,X,T, Mxp, R6, R12, Sx Basic System Configuration guide.
- To know more about the release that supports satellite mode of operation, please refer to the 7x50 SROS 15.0 release notes.

In satellite mode of operation, 7210 SAS-Sx chassis act as a 10G port extender for the SR host and allows the operator to utilize the maximum per slot switching capacity on the host when delivering services over 10GE ports. In this mode, it is expected that the operator will typically connect the 100GE ports of the satellite (referred to as satellite uplinks) to the 100GE ports on the SR host and connect the 10GE ports on the satellite to customers. In the satellite mode of operation, all configuration and management functions are performed through the host node, with the satellite appearing to be a remote MDA hanging of the SR node. There is no option available to manage the node operated in this mode. As such, access to satellite console port, management port and SNMP is not available and no IP address needs to be assigned to the satellite. All the services and service attributes (For example: QoS, ACLs, etc.) for the customers connected to the satellite 10GE ports are provisioned on the SR host. The service functionality seen by the customers is close to the functionality had they been directly connected to the ports of the IMM on the host.

On power-on, with the chassis-role set to factory-default or satellite, the device will automatically initiate request to download the BOF and the Timos image to bootup, using the 100GE uplinks. After successful bootup it will obtain the configuration information required to function as a satellite. CLI commands are available on the host to configure the satellite information so that it can recognize the satellite and provide it with the correct Timos image. For more information on successfully booting up the satellite for the first time refer to the 7210 SAS-Sx 10/100GE Installation guide.

For more information about configuring the satellites, the services supported and other details refer to the 7x50 SROS user guide manuals.

8.1.3 IPv4 unnumbered support with MPLS.

Platforms Supported: 7210 SAS-M network mode, 7210 SAS-T network mode, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-R6 and 7210 SAS-R12.

This release adds support for unnumbered IPv4 interfaces along with support for setting up MPLS RSVP and LDP tunnels with unnumbered IPv4 interfaces. This feature eases the provisioning process when IP/MPLS nodes are added to or remove from a ring of IP/MPLS nodes, without affecting the configuration on the adjacent nodes reducing the time involved and the chances of errors. In addition, it also reduces the number of infrastructure IPv4 addresses in use.

- **Support for Ethernet Unnumbered Interfaces**

Starting with this release, the ability to configure Ethernet unnumbered interfaces has been added to support use of IPv4 unnumbered interfaces with MPLS based services. In previous releases, Ethernet unnumbered interfaces was supported on some 7210 platforms for use with MPLS-TP and this release extends support on those platforms to support MPLS.

Unnumbered Ethernet allows point-to-point interfaces to borrow the address from other interfaces such as system or loopback interfaces. This feature enables unnumbered interfaces for some routing protocols (IS-IS and OSPF). Support for routing is dependent on the respective routing protocol and service. This feature also adds support for both dynamic and static ARP for unnumbered Ethernet interfaces to allow inter working with unnumbered interfaces that may not support dynamic ARP.

An unnumbered interface is an IPv4 capability only used in cases where IPv4 is active (IPv4-only and mixed IPv4/IPv6 environments). When configuring an unnumbered interface, the interface specified for the unnumbered interface (system or other) must have an IPv4 address. Also, the interface type for the unnumbered interface will automatically be point-to-point. Unnumbered Ethernet can be used in IES and VPRN access interfaces, as well as in a network interface. It cannot be used with IP interfaces associated with RVPLS service.

- **Support for RSVP-TE with IPv4 unnumbered interfaces**

This release introduces the use of unnumbered IP interface as a Traffic Engineering (TE) link for the signaling of RSVP P2P LSP and P2MP LSP. The support of unnumbered TE link in IS-IS consists of adding a new sub-TLV of the extended IS reachability TLV, which encodes the Link Local and Link Remote Identifiers as defined in RFC 5307.

The support of unnumbered TE link in OSPF consists of adding a new sub-TLV, which encodes the same Link Local and Link Remote Identifiers in the Link TLV of the TE area opaque LSA and sends the local Identifier in the Link Local Identifier TLV in the TE link local opaque LSA as per RFC 4203. The support of unnumbered TE link in RSVP implements the signaling of unnumbered interfaces in ERO/RRO as per RFC 3477 and the support of IF_ID RSVP_HOP object with a new Ctype as per Section 8.1.1 of RFC 3473. The IPv4 Next/Previous Hop Address field is set to the borrowed IP interface address.

The unnumbered IP is advertised by IS-IS TE and OSPF TE, and CSPF can include them in the computation of a path for a P2P LSP or for the S2L of a P2MP LSP. This feature does not, however, support defining an unnumbered interface for a hop in the path definition of an LSP.

All MPLS features available for numbered IP interfaces are supported, with the following exceptions:

- Configuring a router-id with a value other than system.
- Signaling of an LSP path with an ERO-based loose/strict hop using an unnumbered TE link in the path hop definition.

- Signaling of one-to-one detour LSP over unnumbered interface (one-to-one FRR is not supported).
- Soft preemption of LSP path using unnumbered interface.
- Inter-area LSP.
- Unnumbered RSVP interface registration with BFD.
- RSVP Hello and all Hello related capabilities such as Graceful-restart helper.
- The user SRLG database feature — The user-srlg-db option under MPLS allows the user to manually enter the SRLG membership of any link in the network in a local database at the ingress LER. The user cannot enter an unnumbered interface into this database and as such all unnumbered interfaces will be considered as having no SRLG membership if the user enabled the user-srlg-db option.

This feature also extends the support of lsp-ping, p2mp-lsp-ping, lsp-trace, and p2mp-lsp-trace to P2P and P2MP LSPs that have unnumbered TE links in their path.

- **Support for LDP with IPv4 Unnumbered Interfaces**

This release allows LDP to establish Hello adjacencies and to resolve unicast and multicast FECs over unnumbered LDP interfaces. Hello adjacencies will be brought up using link Hello packets with the source IP address set to the interface-borrowed IP address and a destination IP address set to 224.0.0.2. The transport address for the TCP connection, which is encoded in the Hello packet, will always be set to the LSR-ID of the node. The source and destination IP addresses of LDP packets are the transport addresses (i.e., LDP LSR-IDs) of the LDP peers.

A FEC can be resolved to an unnumbered interface in the same way as it is resolved to a numbered interface. The outgoing interface and next-hop are looked up in RTM cache. The next-hop consists of the router-id and link identifier of the interface at the peer LSR. This feature supports resolving an LDP FEC over ECMP next-hops consisting of a mix of unnumbered and numbered interfaces. All LDP FEC types are supported.

This feature also extends the support of “lsp-ping”, “p2mp-lsp-ping”, and “ldp-trace” to test an LDP unicast or multicast FEC that is resolved over an unnumbered LDP interface.

For more information, please refer to the 7210 SAS M, X,T, Mxp, R6, R12, Sx MPLS Guide, 7210 SAS M,X,T, Mxp, R6, R12, Sx Routing Protocols Guide and 7210 SAS M, T, Mxp, Sx Services Guide.

8.1.4 BFD

Platforms Supported: 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE.

This release adds support BFD on SAS-Sx/S 1/10GE platforms. BFD on these platforms is implemented in hardware with support for 10ms timers. It allows for faster failure detection and triggering of specified actions to improve network and service resiliency. 10ms timers is allowed only for BFD sessions that is configured with an IP interface associated with a port. BFD for IP interface over a LAG or using the system IP interface is implemented in the CPM and supports 100ms timers. BFD support on these platforms is available with the following services and protocols:

- BFD in a VPRN service can be used for:
 - OSPV2 PE-CE routing protocol
 - Static routes (only IPv4)
 - VRRP (IPv4)
 - BGP for PE-CE protocol (IPv4)
 - PIM for PE-CE protocol for NG-MVPN (IPv4)
- BFD in IES service can be used for:
 - OSPFv2
 - IS-IS for IPv4 interfaces
 - Static routes (only IPv4)
 - VRRP (IPv4)
- BFD in Base routing instance can be used for:
 - OSPFv2 on network IPv4 interfaces
 - IS-IS on network IPv4 interfaces
 - VRRP on network IPv4 interfaces
 - MP-BGP for vpn-ipv4 and vpn-ipv6 family (only multi-hop)
 - Static routes (only IPv4)
 - RSVP-TE
 - PIM (IPv4)
 - TLDP (IPv4)
 - Interface LDP (link-level) (IPv4)

Some of the exceptions are - BFD is not supported for IPv6 interfaces in both base routing instance and IPv6 VPRN services.

For more information, please refer to the 7210 SAS M,X,T, Mxp, R6, R12, Sx Router Configuration guide, 7210 SAS M, T, Mxp, Sx Services guide and 7210 SAS M,X,T, Mxp, R6, R12 ,Sx Routing Protocols guide.

8.1.5 PTP with support for IEEE default profile and G.8265.1 profile

Platforms Supported: 7210 SAS-Sx 1/10GE fiber variants (i.e. 7210 SAS-Sx 46F 2C 4SFP+ and 7210 SAS-Sx 22F 2C 4SFP+)

This release supports use of 1588v2/PTP for frequency and time recovery. It supports both OC slave and BC with IEEE default profile and ITU-T G.8265.1 frequency profile. It also supports the use of 1PPS OUT interface (present on the rear of the chassis) for measuring accuracy of the recovered PTP clock.

This release also supports PTP Hybrid mode (that is, with syncE assist for frequency recovery). Additional CLI command “configure>system>ptp>clock>freq-source ssu” is available to enable syncE assist.

NOTES:

- Support for 7210 SAS-Sx 1/10GE copper and PoE variants is not available in this release. This release also does not support 7210 SAS-S 1/10GE platforms.
- 1PPS IN/OUT interface is not supported in this release. For more information, see the 7210 SAS M,X,T, Mxp, R6, R12, Sx Basic System Configuration guide.

8.1.6 PTP – support for G.8275.1 profile and Ethernet encapsulation (BETA only)

Platforms Supported: 7210 SAS-R6 and 7210 SAS-R12.

This release supports IEEE 1588 messaging using direct encapsulation in Ethernet frames, as per Annex F of the standard, along with support for port-based timestamping (PBT). This encapsulation mechanism is configured on a per-port basis. In addition, this release enhances the IEEE 1588 functionality with the addition of support for the ITU-T G.8275.1 profile. This profile specifies how to use IEEE 1588 to distribute a time reference to meet the requirements of mobile base-stations.

NOTE: This is a BETA version of the feature and must not be used in production networks.

For more information, please refer to the 7210 SAS M,X,T, Mxp, R6, R12, Sx Basic System Configuration guide.

8.1.7 Netconf/Yang (BETA only)

Platforms Supported: 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This release introduces the support for NETCONF. NETCONF is an IETF network management protocol published as RFC 6241. It runs on top of the SSHv2 transport protocol (SSHv2 is an existing protocol supported on 7210) as specified in RFC 6242. NETCONF can be used as an alternative to CLI or SNMP for managing the node. It is an XML-based protocol used to configure network devices and uses RPC messaging for communicating between a NETCONF client and the NETCONF server running on the 7210 node. An RPC message and configuration data are encapsulated within an XML document. These XML documents are exchanged between a NETCONF client and a NETCONF server in a request/response type of interaction. The SR OS NETCONF interface supports both configuration and retrieval of operational information. In this release, base capability 1.0 and 1.1 (primarily chunked framing) is supported. The SR OS NETCONF supports the XML format content layer requests and responses for configuration data. Both <edit-config> and <get-config> NETCONF requests support an XML format for the content layer.

NOTE: The NETCONF port is not configurable. NETCONF sessions are supported on TCP port 830 (as required in RFC 6242). NETCONF sessions received on other TCP ports (including 22) are not supported.

This release also provides a full YANG data model for all 7210 SAS configuration data (the equivalent of everything under the CLI configure context). The YANG data model is composed of a set of proprietary Nokia YANG modules that are distributed as part of the SR OS image file bundle. XML formatted configuration data that conforms to the YANG modules is supported by the SR OS NETCONF server.

For more information, please refer to the 7210 SAS System Management User Guide.

8.1.8 IPv4 VPRN service

Platforms Supported: 7210 SAS-K 2F4T6C.

RFC 4364 details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers. Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association. The service provider uses BGP to exchange the routes of a VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers belonging to the same VPN. Since the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm. When BGP distributes a VPN route, it also distributes an MPLS label for that route. On 7210 SAS devices, a single label is assigned to all routes in a VPN. Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label header, so that it gets tunneled across the backbone to the proper PE router.

VPRN services can be used in mobile-backhaul applications to provision routing instances for different mobile service protocols typically involved in mobile backhaul such as OAM, signaling, Voice, Data, etc. It can also be used to provide IP VPN services for business customers. The scaling on the 7210 SAS-K 2F4T6C is designed to fit a hierarchical network design/seamless MPLS design and it is highly recommended to use these design principles when deploying 7210 SAS-K 2F4T6C. TWAMP-light is supported in a VPRN service for service performance monitoring, allowing operators to measure SLAs for service delivery.

In this release, the following support is available:

- L3 SAP types of NULL, Dot1q and QinQ
- Service label per VRF
- Static spoke-sdp for PE to PE
- VRF export/import policies
- QoS – Per service ingress and service egress queuing, hierarchical shaping and scheduling, along with following major functionality:
 - SAP ingress classification using mac-criteria or ip-criteria (5-tuple and DSCP)
 - SAP egress marking - DSCP and Dot1p bits
 - MPLS EXP bits used for network side classification and marking for MPLS LSPs used to transport VPRN service packets

- ACLs/Filter - SAP ingress/egress ACLs (ip-criteria only)
- MP-iBGP protocol (vpn-ipv4 family only) for VPRN route exchange
 - 4-Byte ASN support available (enable/disable)
 - BGP route policy support (export/import policies)
 - Use of Route reflector is highly recommended (with 7210 configured as client only)
 - Use of Outbound Route Filtering (ORF) to retrieve only those routes that are needed by the node to deliver services configured on it
 - VPRN aggregate route support is available
 - PE-CE routing protocol – eBGP (IPv4 family), OSPFv2, and Static routing with CPE-check connectivity.
- Use of RSVP-TE LSPs or LDP LSPs to reach the VPRN route MP-BGP nexthop
- OAM Tools
 - ICMP ping and traceroute applications available in VPRN context
 - vpn-ping and vpn-traceroute OAM tools (SAA support available)

In this release, the following functionality is not available (or available as BETA as noted below):

- Ethernet Spoke-SDP termination
- BGP Multipath/ECMP
- BFD for BGP
- Route leak between VRF, including GRT leaking using Route Policy
- BGP 3107 labeled unicast route support with service optimization is BETA only

For more information, please refer to 7210 SAS-K 2F4T6C User Guides.

8.1.9 OSPFv2 with TE

Platforms Supported: 7210 SAS-K 2F4T6C.

OSPFv2 (Open Shortest Path First version 2) is an IGP and a link state protocol used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

Major OSPF features that are supported on 7210 SAS-K 2F4T6C are:

- Stub areas
- Not-So-Stubby areas (NSSAs)
- Virtual links
- Authentication
- Route redistribution
- Routing interface parameters
- OSPF-TE extensions (supports CSPF and MPLS fast reroute)
- OSPF instances (only a single instance is supported, but instance ID can be different)

7210 SAS-K 2F4T6C IPv4 FIB scaling is designed for it to be deployed in a stub or NSSA area or in a routing domain (using OSPF instance) and it is highly recommended to use these network design principles.

Please refer to the 7210 SAS-K 2F4T6C Routing protocols User Guide for more information.

8.1.10 BGP RFC 3107 Labeled unicast routes for L2 services, along with service optimization

Platforms Supported: 7210 SAS-K 2F4T6C.

This release adds support for use of inter-AS or intra-AS BGP RFC 3107 labeled unicast routes for transport of L2 VPN services (Epipe and VPLS). This allows 7210 SAS-K 2F4T6C to be used in a seamless MPLS network using BGP 3107 labeled unicast routes to transport services.

A new type of transport method is supported under SDP that allows using BGP LSP tunnels to reach the far end. The 7210 SAS devices support the 3-label version of BGP LSP tunnels. On 7210 SAS-K 2F4T6C, both FRR facility and FRR one-to-one is supported for RSVP LSPs used with RFC 3107 BGP LSPs. In this release, only Option-C is supported.

In a seamless MPLS deployment, using BGP 3107 labeled routes for transport tunnels, access nodes can potentially receive 1000s of addresses for service endpoints (PEs) corresponding to all the nodes in the network. With the cost-effective scale of 7210 SAS-K 2F4T6C, it is not possible to accommodate large number of labeled unicast routes into the MPLS FIB, though a larger number of routes can be held in the BGP RIB. 7210 provides an option to install only those BGP labeled

routes that are needed for services configured on the node, alleviating the need to install all the received routes and making it possible to use smaller MPLS FIB. This option is called 'BGP 3107 labeled route service optimization'. When this option is enabled, BGP 3107 transport tunnel routes are held in the RIB and one of the following triggers installation of the routes into the MPLS FIB:

- SDP configuration to far-end destination (for Epipe and VPLS services configured on the node), with the option to use BGP 3107 tunnels enabled
- Dynamic SDP-binding (PW) creation when BGP-AD routes are received for an VPLS service configured on the node (future release)
- L3 VPN routes that are reachable only through BGP 3107 labeled routes (BETA only)

NOTES:

- The BGP 3107 service optimization option is recommended for use only on the edge PE devices. When this option is enabled, IPv4 labeled routes not required for use by services (for example management IPv4 routes not used in services) are not installed in the FIB.
- BGP for exchange of IPv4 non-labeled routes is not supported on 7210 SAS-K 2F4T6C. It must not be used. Only RFC 3107 labeled routes for L2 services is supported and validated. In software, it has not been blocked though. Additionally, use of BGP labeled route is supported only for /32 prefixes.
- 7210 SAS-K 2F4T6C does not support the ASBR functionality for BGP 3107 routes.
- In this release, BGP 3107 labelled route nexthop is resolved using RSVP LSPs only. Use of LDP LSPs for nexthop resolution is not supported in this release.

For more information refer to the 7210 SAS-K 2F4T6C Routing Protocols User Guide, 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Services User Guide.

8.1.11 LSP ping and trace extensions

Platforms Supported: 7210 SAS-K 2F4T6C.

With this release, the coverage of the LSP ping and trace tools, is extended to test connectivity of an LSP using a BGP RFC 3107 label route. Support of the target FEC stack TLV of type BGP Labeled IPv4 /32 Prefix as defined in RFC 4379 has also been added.

Note that only BGP label IPv4 /32 prefixes are supported because these are usable as tunnels in 7210. BGP label IPv6 /128 prefixes are not currently usable as tunnels on the 7210 platform and as such, are not supported in LSP ping/trace.

For more information, please read the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OAM User Guide.

8.1.12 Y.1564 for 7210 SAS-K 2F4T6C

Platforms Supported: 7210 SAS-K 2F4T6C.

ITU-T Y.1564 defines the test methodology to be used and parameters to be measured to test service SLA conformance during service turn up. It primarily defines 2 test phases. The first test phase defines service configuration test, which consists of validating whether the service is configured properly. As part of this test the throughput, Frame Delay, Frame Delay Variation (FDV), and Frame Loss Ratio (FLR) is measured for each service. This test is typically run for a short duration. The second test phase consists of validating the quality of services delivered to the end customer and is referred to as the service performance test. These tests are typically run for a longer duration and all traffic is generated up to the configured CIR for all the services simultaneously and the service performance parameters are measured for each the service.

In this release, 7210 SAS-K 2F4T6C supports service configuration test for user configured CIR rate and measurement of delay, delay variation and loss with the testhead OAM tool. 7210 testhead OAM tool supports bi-directional measurement and it can generate test traffic for only one service at a given time. Testhead OAM tool can be used to perform out-of-service tests for VPLS and Epipe services. It can validate if the user specified rate is available and compute the delay, delay variation and loss for the service under test at the specified rate.

7210 SAS-K 2F4T6C can generate traffic for up to four streams simultaneously allowing users to validate single FCs of multiple SAPs or multiple FC of a single SAP or a mix-n-match of both. It can generate traffic up to about 1Gbps rate across all the streams at a given time.

Additionally, per SAP loopback with mac-swap is supported in this release. It must be used at both ends before using the testhead OAM tool. It allows user to perform Y.1564 based service activation tests without disrupting other services configured on the same port as the test SAP.

For more information, please read the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OAM User Guide.

8.1.13 G.8032

Platforms Supported: 7210 SAS-K 2F4T6C.

ITU-T G.8032 specification defines protocol mechanisms to provide Ethernet ring protection to enable deployment of resilient Ethernet Layer 2 networks. G.8032 (ETH-ring) is built on Ethernet OAM and is also referred to as Ring Automatic Protection Switching (R-APS).

Eth-rings are supported for protection of only VPLS services (including routed VPLS) and can be used with both access SAPs and access uplink SAPs configured in the service. ETH-ring multi-homing into other ETH-rings or VPLS PEs is also supported in the current release. ETH-rings offer fast resiliency for Ethernet services leveraging ring topologies for any single link or node failure. By configuring multiple ETH-rings instances on the same physical topology, G.8032 can utilize all link resources in a ring by load-balancing the service traffic over the two arms of the ring.

7210 SAS-K 2F4T6C can be part of both major rings and sub-rings. It cannot be configured as an inter-connection node used to connect major and sub-rings together. It supports use of CCMs with 10ms timers for failure detection, along with support for fast-flood mechanisms for faster service restoration on failures in the ring.

For more information about this feature, refer to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Interfaces User Guide and 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Services User Guide.

8.1.14 OAM-PM

Platforms Supported: 7210 SAS-K 2F4T6C.

OAM tools used for measuring and monitoring performance metrics such as delay, inter-frame delay, loss and availability for Layer-2 VPN services and Layer-3 services has been enhanced significantly with the capability included for collecting and storing results in fixed time-based measurement intervals. Option is available to store results in fixed-time-based measurement intervals (5 min, 15min, 1hr and 24hr). Delay measurement (Frame Delay, Inter-Frame Delay and Frame Delay Range) results will be stored in bins that are defined in microseconds by a lower threshold and an upper limit and the Mean Frame Delay is computed over the length of the measurement interval. Loss measurement computes frame loss ratio over fixed time-intervals and

is used to monitor service availability by computing high-loss intervals and consecutive high-loss intervals. In addition, Threshold Crossing Alarms (TCA) can be configured against the various key performance metrics collected as part of the OAM-PM architecture. Both stateful and stateless TCA are supported for delay, loss and availability metrics.

For Layer-2 services, service OAM performance-monitoring tools using ITU-T Y.1731 DMMv1 and SLM can be used. DMMv1 messages are used to measure and monitor delay, Inter-frame delay and frame delay range and SLM messages are used for loss and availability. For Layer-3 services – VPRN and IES services, IETF TWAMP-light can be used for monitoring delay and loss.

For more information, please refer to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OAM User Guide.

8.1.15 OAM - TWAMP-light

Platforms Supported: 7210 SAS-K 2F4T6C.

TWAMP-light support has been added to the OAM-PM architecture for the collection and reporting of delay, loss and availability metrics. TWAMP Light uses the standard TWAMP packet but removes the TCP control channel using local configuration on the session controller and session reflector. Support is available for test launching and reflection in both the base router instance and on a per-VRPN basis. It allows operators to measure service SLA metrics such as delay, jitter and loss for IP services using the OAM-PM infrastructure and deliver performance assured services to their customers.

On 7210 SAS-K 2F4T6C, TWAMP-light is supported in base routing and VPRN service instance with IPv4 interfaces.

For more information, please refer to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OAM User Guide.

8.1.16 LDP ECMP

Platforms Supported: 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE.

This release includes support for LDP ECMP for LSR LSPs only. It provides an option of load-balancing traffic across multiple LDP LSPs at an LSR node. This feature allows operator to improve the capacity available for transport of services and network resiliency by adding links which are advertised as LDP ECMP paths.

NOTE: LDP ECMP is not supported on a LER node.

The ECMP resources in hardware are shared with IP ECMP feature. Before using LDP ECMP at a LSR node, user needs to allocate ECMP resources and enable it using the CLI command `configure> system> resource-profile> router> ldp-ecmp-routes`. For more information about resource usage, see the 7210 SAS-M, Sx,T,X,R6, R12, Sx Basic System Configuration User Guide.

NOTE: LDP software will install ECMP paths if hardware resources are available for the /32 prefixes. Software will install only a single path after it runs out of ECMP resources in hardware, subject to the system limit for number of LDP /32 prefixes.

For more information about LDP ECMP, see the 7210 SAS-K 2F4T6C MPLS User Guide.

8.1.17 7210 SAS-K 2F4T6C Feature Support

Platforms Supported: 7210 SAS-K 2F4T6C

The following features which were BETA in the prior release is now available for production deployments:

- IGMP snooping and MVR. 7210 SAS-K 2F4T6C support IGMP v1, v2, and v3 snooping in VPLS services with L2 uplinks and MPLS uplinks.
- CFM (with Down MEP, UP MEP and MIP support) and Y.1731 Ethernet OAM tools
- MPLS OAM tools such as `lsp-ping`, `lsp-trace`, `vccv-trace`, `vccv-ping`
- Y.1564 Testhead OAM tool and per SAP loopback with `mac-swap`

For more information, please refer to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Services Guide and 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OAM and Diagnostics User Guide.

8.2 Release 9.0R3

The following items describe features added to 9.0R3 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

8.2.1 7210 SAS Platform - 7210 SAS-K 2F4T6C (12-port IP/MPLS device) Non-ETR

This release includes support for the following 7210 SAS-K 2F4T6C Non-ETR platform. It is an 1RU platform, providing a total of 12 ports - 2 x 100/1000 SFP ports, 4 x 10/100/1000Mbps RJ-45 Base-T ports, and 6 Combo ports (Base-T or SFP). Along with the combo ports, user can use a maximum of up to 8 x 100/1000 SFP ports or a maximum of up to 10 x 10/100/1000Mbps RJ-45 Base-T fixed copper ports or a mix-n-match of SFP and copper ports, with a maximum of up to 12 Ethernet ports.

It is an IP/MPLS capable device, designed for delivering differentiated services for mobile backhaul, business services and vertical industry (e.g. railways, transportation, etc.). It supports both IP/MPLS and Ethernet based Layer-2 VPN and Layer-3 services, with hierarchical per service ingress and egress queuing, shaping & scheduling with 64Mbytes of packet buffers, which allows for delivering differentiated services with better SLAs to customers. The scale is designed to be sufficient to fit into a seamless MPLS network design and or hierarchical MPLS design. It supports a rich set of Ethernet and IP/MPLS OAM tool set, to aid in fault resolution and service performance monitoring. Synchronous Ethernet (SyncE) and PTP/1588v2 is supported to distribute frequency and time.

The following are some of the hardware features:

- 1RU platform designed to be NEBS compliant
- Front access to all the 1GE service interfaces
- 9Gbps (full-duplex) switching for IMIX traffic
- External power supply (AC) with a single power feed
- Passive cooling (No fans)
- Standard operating temperature range of 0 degree Celsius to 40 degrees Celsius
- 2 storage locations - External field-replaceable SD card (cf1:\), and USB (uf1:\) storage device. SD card slot access is on the rear of the chassis

The following functionality (only major ones listed) is supported with this release:

- Support Access, Access Uplink and Network port mode
- Support the following SAP encapsulations - NULL, Dot1q, QinQ
- Service Support
 - Layer-2 VPN services - Epipe, VPLS
 - Layer-3 services - IES
 - T-PE functionality is available for Epipe service

-
- Epipe service
 - VPLS service with spoke-SDP and mesh-SDP
 - L2PT and BPDU tunneling in VPLS service
 - IGMP (v1, v2, v3) snooping (Layer-2 multicast) and MVR in VPLS service (both SAPs & SDPs) - BETA only
 - DHCP snooping in VPLS services (both SDPs and SAPs)
 - Layer-2 control protocol tunneling support for EFM, LLDP, 802.1x, and LACP
 - IES IPv4 services with support for OSPFv2, IS-IS (IPv4), static routing
 - MPLS Support
 - MPLS support for Epipe and VPLS services
 - MPLS label push and pop supported
 - RSVP-TE with FRR (one-to-one and facility with PHP), primary & secondary LSPs with hot standby, SRLG, admin-groups, etc.
 - IPv4 Routing Support;
 - IPv4 forwarding support with static routing, OSPFv2 and IS-IS (IPv4) routing protocols
 - Route policies for management and control of distribution of routing information is available
 - DHCP server (IPv4) and DHCP relay (IPv4) support for IES services
 - BFD is not supported in this release
 - IPv6 is not supported in this release
 - QoS and ACL Support
 - 8 Forwarding classes (FCs)
 - Per service Ingress and service egress hierarchical QoS with 64Mbytes of buffers
 - SAP ingress QoS with queuing (with 8 queues per SAP), scheduling, and hierarchical shaping (2 levels - per FC, per SAP)
 - SAP egress QoS with queuing (with 8 queues per SAP), scheduling, and hierarchical shaping (2 levels - per FC, per SAP and a per port egress rate shaper)
 - Network port and Access uplink port ingress and egress queuing (with 8 queues per port), scheduling & shaping (2 levels - per FC, per port egress rate shaper)
 - Strict-priority scheduling (SP), WFQ or hybrid (i.e. a mix where some queues are configured in SP and some other queues are configured in WFQ) is supported on SAP ingress, SAP egress, Network ingress and Network Egress

-
- SAP ingress classification supports Dot1p, IP DSCP, MAC criteria, IPv4 criteria, and IPv6 criteria (as applicable)
 - Network port ingress supports Dot1p & DEI, IP DSCP and EXP based classification
 - Access uplink port ingress supports Dot1p & DEI, IP DSCP based classification
 - Egress marking using Dot1p, IPv4 DSCP and MPLS EXP (as applicable)
 - SAP ingress and egress ACLs, Network port IP interface ingress and egress ACLs are supported, with MAC criteria, IPv4 criteria and IPv6 criteria (as applicable). IPv6 match criteria supported only for Epipe and VPLS services in both QoS classification and ACLs
 - Network Synchronization support
 - Synchronous Ethernet is available on all ports (including copper ports)
 - PTP/1588v2 OC slave and BC, with IP/UDP encapsulation is available
 - PTP/588v2 PTP hybrid mode (SyncE for frequency & PTP for time) for OC slave and BC is available
 - PTP/1588v2 profiles - IEEE default profile & G.8265.1 profile is available;
 - High Availability and Reliability Support
 - LAG with active/active and active/standby support
 - MC-LAG support (only client support)
 - G8032 is not supported in this release
 - STP, RSTP, MSTP with mVPLS/xSTP support
 - MPLS FRR - facility with & without PHP and one-to-one support;
 - MPLS primary and secondary LSPs, with hot-standby secondary LSP support
 - Active/Standby PW in Epipe and VPLS services
 - Fault propagation support in Epipe service (For example: LLF, and others.)
 - OAM support
 - EFM OAM with support for EFM OAM dying gasp message or a SNMP dying gasp message on loss of power
 - LLDP - for topology discovery
 - CFM/Y.1731 Down MEP, UP MEP, Ingress/Egress MIP for VPLS and Epipe services (see user guide for MEP support per service and different service objects) - BETA only
 - CFM and Y.1731 based OAM tools - CCM, Linktrace, Loopback, 2-DM, 1-DM, 2-SLM, AIS, and RDI
 - SAP loopback without and with MAC swap is supported

- Y.1564 testhead OAM tool for service performance measurement before service activation is available - BETA only
- MPLS OAM tools for Epipe and VPLS services (example - lsp-ping, vccv-ping, vccv-trace, mac-ping, and others.) - BETA only
- SAA support for selected set of OAM tools for measurement and collection of service performance metrics
- PTP can be used to driver system time, allowing for better accuracies with performance measurement tools used for delay measurement
- Mirroring support - Local Mirroring is supported with port and Dot1q SAP as mirror destination
- Accounting, Security and Management support
 - Per SAP ingress and egress accounting records
 - Per network port accounting records
 - Dot1x is not supported in this release
 - SNMP (including v3 support), SSH, Telnet
 - NTP and SNTP is available
 - RADIUS and TACACS+ supported
 - User profiles are supported
 - Software defines a policy which is used for CPU protection and it is not user configurable
 - Supports Autoinit, which allows operators to deploy the nodes faster. It provides the following boot options:
 - Using Autoinit
 - Using images on the external SD card (cf1:\)
 - Using images on the USB (uf1:\)



Note: The device does not ship with any software images when shipped from factory. To boot the device, software license must be purchased to obtain SD card with software images. Before booting the device, the SD card must be inserted. By default, the device is configured to use autoinit. For more information reference on how to use the various options to boot the device to the 7210 SAS-K Installation Guides and 7210 SAS Basics System Configuration Guide.

For more information refer to the all the 7210 SAS-D, E, K User Guides.

8.2.2 NG-MVPN - RSVP signaling for P2MP LSPs

Platforms Supported - 7210 SAS-Mxp, 7210 SAS-T, 7210 SAS-R6, 7210 SAS-R12.

Prior release added support for Next-Generation Multicast VPNs (NG-MVPN) with BGP control plane and mLDP support for P2MP LSPs. This release includes support for RSVP signaling for P2MP LSPs, providing an option to use RSVP based signaling to users who use RSVP in their networks.

With NG-MVPN support 7210 SAS device can be used to deliver multicast services efficiently to customers in VPRN services, extending the delivery of IPv4 multicast traffic closer to end user in business and mobile backhaul networks.

With this release, Next generation multicast-VPN (mVPN) implementation supports RSVP P2MP LSP for use with I-PMSI and S-PMSI multicast distribution trees. RSVP P2MP must be enabled in the global instance and I-PMSI configuration option must be enabled to allow per mVPN instance to automatically setup P2MP LSP to leaf PE nodes learned through BGP multicast Auto-Discovery signaling. S-PMSI is used to avoid sending traffic to a PE node that participates in the multicast VPN but that does not have any receivers for the multicast flow. It allows for the efficient distribution of multicast traffic over a provider network, specifically for high bandwidth multicast flows. On 7210, spawning of S-PMSI tunnel based on data-threshold is not supported. Instead, operator has an option to configure the use of S-PMSI tunnel for the configured range of multicast groups. In other words, an S-PMSI tunnel is spawned dynamically for the configured multicast flows irrespective of the rate of multicast traffic.



Note: Before using this feature, resources must be allocated for use by this feature from the sf-ingress-internal-tcam pool of resources, using the CLI command `configure> system> global-res-profile> sf-ingress-internal-tcam> mpls-p2mp`. If the 7210 SAS-R is deployed as a bud router, then the CLI command `configure> system> loopback-no-svc-port> p2mpbud <p2mpbud-port-id>`, must be used to configure one of the front-panel ports as a loopback port for use with this feature.

For more information, see the 7210 SAS-X, R6,R12 Services User Guide, 7210 SAS-M, T, Mxp, X, R6, R12, Sx Routing Protocols User Guide, 7210 SAS-M, T, Mxp, X, R6, R12, Sx Basic System Configuration User Guide and 7210 SAS-M, T, Mxp, X, R6, R12, Sx MPLS User Guide.

8.3 Release 9.0R2

There are no new features added to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS release 9.0R2.

8.4 Release 9.0R1

The following items describe features added to 9.0R1 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS.

8.4.1 7210 SAS Platform - 7210 SAS-Sx 1/10GE Copper Variant

This release includes support for the following 7210 SAS-Sx 1/10GE variants - 7210 SAS-Sx 48T 4SFP+ and 7210 SAS-Sx 24T 4SFP+. These are 1RU platform, providing an option of using either 48 x 10/100/1000Mbps Base-Tx ports or 24 x 10/100/1000Mbps Base-Tx ports. In addition it has 4 x 10GE SFP+ ports.

Both these variants provide the following support (unless specified otherwise):

- 1RU platform designed to be NEBS compliant
- Front access to all the 1GE and 10GE service interfaces
- Line-rate switching on all the ports
- Redundant hot-swappable power supply (AC, DC -48V)
- Non-removable Fans on the rear of the chassis
- 3 storage locations – Internal non-replaceable flash (cf1:\) of size 2GB, external field-replaceable SD card (cf2:\), and USB (uf1:\) storage device. USB and SD card slot access is on the rear of the chassis
- Console port and Out-of-band Ethernet management interface on the rear of the chassis

**Note:**

This release also supports the other 7210 SAS-Sx 1/10GE variants 7210 SAS-Sx 46F 2C 4SFP+, 7210 SAS-Sx 22F 2C 4SFP+, 7210 SAS-Sx 48Tp 4SFP+ PoE, and 7210 SAS-Sx 24Tp 4SFP+ POE.

The platforms can be operated in two modes of operations:

- Satellite mode of operation
- Standalone mode of operation

**Note:**

- User has an option to configure the role of the chassis using the BOF menu to configure the BOF parameter 'chassis-role' and 'host-type'. The BOF parameter 'chassis-role' allows the user to nail down the chassis to operate in either satellite mode or standalone mode of operation. For more details on use of this parameter use the 7210 SAS-Sx Basic Systems user guide.
- To know more about the release that supports satellite mode of operation, please refer to the 7x50 release notes.
- For more information refer on satellite mode of operation and standalone mode of operation to the all the 7210 SAS-M, T, Mxp, Sx, X, R6, R12 User Guides.

8.4.2 IPv6 support (Including 6VPE and 6PE)

Platforms Supported – 7210 SAS-Sx 1/10GE (standalone mode).

Release 9.0R1 adds IPv6 support on 7210 SAS-Sx 1/10GE operating in standalone mode. It can be used as dual-stack IPv6 and IPv4 routers capable of IPv6 forwarding and providing IPv6 services, including IPv6 VPN (6VPE) services. In addition, 6PE is supported to allow inter-connection of IPv6 networks using an IP/MPLS infrastructure. IPv6 support can also be used for management of the node (both in-band and out-of-band support is available).

The following support is available:

- IPv6 support in the base routing instance, with IPv6 network IP interfaces with support for static routing, OSPFv3 and IS-IS for IPv6.
- IPv6 IES interfaces for services with support for static routing, OSPFv3 and IS-IS for IPv6.

- IPv6 VPRN service (6VPE) for connecting IPv6 customer networks. For PE-CE routing, static routing, and PE-CE EBGP is supported.
- 6PE support in the base routing instance with BGP 3017 labeled routes to interconnect IPv6 networks over IP/MPLS infrastructure.
- IPv6 support in the base routing instance for management with support for SSHv6 and other management protocols such as SNMP, Telnet, etc.
- TWAMP-light is supported for performance measurement with IPv6 IP interface in the base routing instance and in VPRN services.
- IPv6 support with out-of-band management Ethernet port.

Some of the restrictions in this release are:

- IPv6 is not supported for use with autoboot/autoinit.
- VRRP and BFD for IPv6 interfaces is not supported.
- NTP and TWAMP server is not supported with IPv6.
- IPv6 multi-cast is not supported.
- IPv6 ECMP is not supported.



Note: Before using IPv6, users need to allocate routing FIB resources for IPv6 using the command `configure> system> resource-profile> route> max-ipv6-routes` and `configure> system> resource-profile> route> max-ip-subnets`. Please see the 7210 SAS-M, T, X, Sx, Mxp Basic System Configuration guide for more information.

For more information, please refer to 7210 SAS-X,R6,R12 Services User guide, 7210 SAS-M,T,X,Mxp,Sx,R6,R12 MPLS User guide, 7210 SAS-M,T,X,Mxp,Sx,R6,R12 Routing Protocols User Guide, 7210 SAS-M,T,X,Mxp,Sx,R6,R12 System Management User guide, and 7210 SAS-M,T,X,Mxp,Sx,R6,R12 Basic System Configuration User Guide.

8.4.3 IPv4 ECMP in the base routing instance

Platforms Supported: 7210 SAS-Sx 1/10GE standalone mode.

With release 9.0R1, support for ECMP (Equal Cost Multiple Path) is available for IPv4 routes. It is supported only in the base routing instance. With this feature, operators can use multiple IPv4 routes to the same destination to increase the bandwidth available for services and to improve service reliability. The IPv4 traffic is load-balanced over the available IPv4 routes to the destination. A maximum of up to 16 routes per destination can be used. IPv4 ECMP is supported for routes configured statically, and learnt through either OSPFv2 or IS-IS (v4).



Note: In addition to enabling use of ECMP by using the command *"configure> router> max-ecmp-routes"*, before using the feature, user must configure the maximum number of IPv4 ECMP routes they plan to use using the CLI command *"configure> system> resource-profile> router> ecmp <max-routes-dest >"*

Note that *ecmp <max-routes-dest >* configuration takes effect only after node re-boot of node. Therefore, it is required to save configuration and re-boot the node. This command determines the maximum number of ECMP routes allowed to be used. Software uses it to allocate ECMP resources in groups and uses the same group for all those ECMP routes that use the exact same set of next-hops to reach the destination. More information of this CLI command is available in the 7210-M, T, X, Mxp, Sx, R6, R12 Basic Systems User Guide.

ECMP is not supported by 7210 for LDP (both LER and LSR), LDPoRSVP, BGP 3107 IPv4 Labelled routes, VPRNs (both IPv4 and IPv6) and IPv6 routes.

Some of the restrictions that apply to this feature are:

- LDP or LDPoRSVP is not recommended for use when using IPv4 ECMP.
- IPv6 routing and interfaces are not recommended for use when using IPv4 ECMP.

For more information refer to the following user guides - 7210 SAS-M,T,X,Mxp,Sx,R6,R12 Routing Protocols User Guide, 7210 SAS-M,T,X,Mxp,Sx,R6,R12 System Management User guide, and 7210 SAS-M,T,X,Mxp,Sx,R6,R12 Basic System Configuration User Guide.

8.4.4 IPv4 Multicast with PIM-SM and IGMPv3 in Base Routing Instance

Platforms Supported: 7210 SAS-Sx 1/10GE (standalone mode).

With this release support is available for IPv4 multicast with support for PIM and IGMPv3. It allows for efficient distribution of multicast traffic in the access networks. 7210 SAS devices support PIM-SM and PIM-SSM along with IGMPv3 in the base routing instance. Following is the list of supported functionality:

- PIM SM (Sparse Mode) and PIM SSM (Source Specific Multicast) for IPv4 multicast.
- IGMP v1, v2, and v3 support with SSM translate.
- PIM DR configuration, with support for use of BFD for detection of DR failure in a redundant configuration.
- Only unicast routing table is used for RPF checks.

- PIM and IGMP route policies can be used to filter join messages.
- PIM RP support – 7210 SAS devices support RP discovery through Static RP configuration, Dynamic RP discovery using BSR protocol and Anycast RP discovery. Support is also available for 7210 SAS device to be used as a RP or as a BSR.
- Static multicast configuration is supported.



Note: SAP ingress QoS policies for IES interfaces and network port ingress QoS are enhanced to classify multicast traffic and control the amount of traffic accepted. Multicast traffic can be classified only when PIM is enabled on the IES IP interface or the network port IP interface.

Following are some of the restrictions:

- On 7210 SAS devices, on ingress of a port multicast traffic can be processed in the context of either igmp-snooping (L2 multicast) or I3-multicast, but not both. In other words, it is not possible to configure SAPs on the port, such that one SAP is a receiver for multicast traffic to be processed by IGMP snooping and another is receiver for multicast traffic to be processed by IP/L3 multicast. An option per port will be available to enable one or the other. By default, IGMP snooping is enabled to be backward compatible. User needs to explicitly change this to allow processing of received multicast traffic by IP/L3-multicast.
- If a VPLS SAP is configured on the same port as the port on which IP multicast is enabled, then multicast traffic received on the SAP is dropped. Unicast, Broadcast and unknown-unicast packets received on the SAP are forwarded appropriately. This behavior is true only for VPLS SAPs and does not apply to VPLS SDPs, Epipe SAPs and Epipe SDPs.

Refer to the 7210 SAS-M, T, X, Mxp, Sx, R6, R12 Routing Protocols User guide, 7210 SAS-M, T, Sx, Mxp Services guide, and 7210 SAS-M, T, Sx, Mxp QoS guide for more information.

8.4.5 TWAMP-light

Platforms Supported – 7210 SAS-R12, 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-T (network mode and access-uplink mode), and 7210 SAS-M (network mode and access-uplink mode), 7210 SAS-K.

In Release 9.0R1, OAM TWAMP Light has been added to the OAM-PM architecture for the collection and reporting of delay, loss and availability metrics. TWAMP Light uses the standard TWAMP packet but removes the TCP control channel using local configuration on the session controller and session reflector. Support is available for test launching and reflection in both the base router instance and on a per-VRPN basis, with both IPv4 and IPv6 interfaces.



Note: In access-uplink mode, TWAMP-light is supported only in base routing instance with IPv6 interfaces configured on access-uplink SAPs.

Refer the 7210 SAS-M, T, X, Mxp, Sx, R6, R12 OAM User guide and 7210 SAS-E,D,K OAM User guide for more information.

8.4.6 OAM-Performance Measurement (OAM-PM)

Platforms Supported – 7210 SAS-R12, 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-T (network mode and access-uplink mode), and 7210 SAS-M (network mode and access-uplink mode), 7210 SAS-K.

In Release 9.0R1, OAM tools used for measuring and monitoring performance metrics such as delay, inter-frame delay, loss and availability for Layer-2 VPN services and Layer-3 services has been enhanced significantly with the capability included for collecting and storing results in fixed time-based measurement intervals. Option is available to store results in fixed-time-based measurement intervals (5 min, 15min, 1hr and 24hr). Delay measurement (Frame Delay, InterFrame Delay and Frame Delay Range) results will be stored in bins that are defined in microseconds by a lower threshold and an upper limit and the Mean Frame Delay is computed over the length of the measurement interval. Loss measurement computes frame loss ratio over fixed time-intervals and is used to monitor service availability by computing high-loss intervals and consecutive high-loss intervals.

For Layer-2 services, service OAM performance-monitoring tools DMMv1 and SLM can be used. DMMv1 messages are used to measure and monitor delay, Inter-frame delay and frame delay range and SLM messages are used for loss and availability. For Layer-3 services – VRPN and IES services, TWAMP-light can be used for monitoring delay and loss.

Refer to the 7210 SAS-M, T, X, Mxp, Sx, R6, R12 OAM User guide and 7210 SAS-E,D,K OAM User guide for more information.

8.4.7 LFA (Loop Free Alternate) SFP policies for LDP FRR

Platforms Supported – 7210 SAS-R12, 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-T (network mode), and 7210 SAS-M (network mode).

The LFA SPF policy feature is supported with both OSPFv2 and IS-IS (IPv4) routing protocols. It is used to enable policy control for computation of backup paths used with LDP FRR.

This feature provides the user with policy control of the Loop-Free Alternate (LFA) backup next-hop selection within Shortest Path First (SPF) calculation in IS-IS (IPv4) and OSPFv2. It introduces the concept of route next-hop template to influence LFA backup next-hop selection. The template supports the following policy attributes:

- New IP Admin-Group include/exclude constraints.
- New IP Shared Risk Loss Group (SRLG) constraints.
- Protection type preference: link or node protection.
- Next-hop type preference: IP (Tunnel next-hop type is not supported on 7210).

The route next-hop template is applied to an IPv4 interface. All prefixes resolved to the interface as primary next-hop have their LFA next-hop selected after applying the policy. This feature supports only network IP interfaces. It also allows the user to exclude prefixes from LFA next-hop calculation by applying an exclude statement within an IGP instance applied to a prefix policy.



Note: Only IP Loop Free Alternate (LFA) is supported for IPv4 routes for use with LDP FRR. IP FRR is not supported for both IPv4 and IPv6 routes. LFA SFP policies are not supported for IES and VPRN IP interfaces.

Please read the 7210 SAS M,T,X,Mxp,Sx,R6,R12 Routing Protocols user guide and 7210 SAS M,T,X,Mxp,Sx,R6,R12 MPLS user guide for more information.

8.4.8 LSP ping/trace for an LSP using BGP 3107 label route

Platforms Supported – 7210 SAS-R12, 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-T (network mode), and 7210 SAS-M (network mode).

With release 9.0R1, the coverage of the LSP ping and trace tools, is extended to test connectivity of an LSP using a BGP RFC 3107 label route. Support of the target FEC stack TLV of type BGP Labeled IPv4 /32 Prefix as defined in RFC 4379 has also been added.

Note that only BGP label IPv4 /32 prefixes are supported because these are usable as tunnels in 7210. BGP label IPv6 /128 prefixes are not currently usable as tunnels on the 7210 platform and as such, are not supported in LSP ping/trace.

Please read the 7210 SAS M,T,X,Mxp,Sx,R6,R12 OAM and Diagnostic User Guide for more information.

8.4.9 IPv4 Secondary address

Platforms Supported – 7210 SAS-R12, 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-T (network mode), and 7210 SAS-M (network mode).

With release 9.0R1, user is provided an option to assign multiple IPv4 addresses to an IPv4 interface. The additional IPv4 addresses assigned to the interface are termed secondary IPv4 addresses. IPv4 secondary addresses allow for use of single VLAN or a LAN for connecting routers/hosts using multiple IP sub-network addresses. This is useful in IP address migration for networks and or use of single VLAN for multiple IP subnets.

IPv4 secondary address assignment is supported with IES, VPRN and network IP interfaces. It can be used for addressing in IP packets and is advertised through OSPFv2 and IS-IS routing protocols. It can also be used in static routing and be offered protection using VRRP when used as a default gateway router address by hosts in the same subnet.

Following are some of the restrictions that apply to use of IPv4 secondary address:

- It is not supported for use in setting up OSPF or IS-IS neighbor-ship.
- It is not supported for setting up MPLS LSPs or PWs.



Note: 7210 supports use of multiple IPv6 addresses with an IPv6 interface on all platforms that support IPv6.

Please read the 7210 SAS M,T,X,Mxp,Sx,R6,R12 Router Configuration user guide for more information.

8.4.10 Percentage rate option for access egress port queues

Platforms Supported: 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Mxp.

With this release, option is provided to configure the access egress port queue CIR and PIR rate as percentages. QoS access-egress policy is used when port-based queuing, shaping and scheduling is enabled.

The percentage rate is relative to the port bandwidth, in other words, the actual value used is dependent on the line-rate speed the port operates at or the limit enforced using the port egress-rate command. In other words, if the port is a 10GE port, the percentages are relative to the 10Gbps speed and on a 1GE port, it is relative to 1Gbps speed. Similarly is egress-rate command is applied to limit the port speed to 100Mbps, on 1GE port, the rate is computed as a percentage of 100Mbps.

It simplifies operations by allowing users to configure a single policy or smaller number of policies that apportions bandwidth to different forwarding class as a percentage of available bandwidth and then apply these policies to different ports in different portions of the network without the need to define multiple different policies with queue rates configured to account for bandwidth.

Please read the 7210 SAS M, T, Mxp, Sx QoS User Guide and 7210 SAS X, R6, R12 QoS User Guide for more information.

8.4.11 Support for IP DSCP marking on access and access-uplink ports

Platforms Supported: 7210 SAS-M access-uplink mode and 7210 SAS-T access-uplink mode.

With this release, option is available for users to enable marking of IP DSCP values for L2 service packets (that is, packets processed in the context of Epipe and VPLS services) sent out of access port or access uplink port. User has an option to configure either Dot1p marking or IP DSCP marking or both.



Note: IP DSCP marking is performed only for IPv4 and IPv6 packets. In addition, the number of VLAN tags in the packet received must match the number of SAP tags to which it is mapped to. If there are more number of tags in the packet, then the IP DSCP value is not modified.

Refer to the 7210 SAS-M, T, Mxp, Sx QoS User Guide for more information.

9 Enhancements

9.1 Release 9.0R4

There are no new enhancements added to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS release 9.0R4.

9.2 Release 9.0R3

There are no new enhancements added to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS release 9.0R3.

9.3 Release 9.0R2

There are no new enhancements added to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS release 9.0R2.

9.4 Release 9.0R1

The following items describe enhancements added to the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-Sx OS Release 9.0R1.

9.4.1 PTP scaling improvement

Platforms Supported – 7210 SAS-M (network mode and access-uplink mode).

With release 9.0R1, the number of PTP slaves that can talk to 7210 SAS-M acting as a boundary clock (BC) has been improved. It is highly recommended to use syncE along with PTP, with syncE providing frequency reference and PTP providing time reference. The amount of CPU processing cycles available influences the amount of PTP slaves that can be connected to the node.

Refer to the 7210 SAS-M,T,X,Mxp,Sx,R6,R12 Basic System User guide for more information.

9.4.2 BFD support for static route in VPRN service

Platforms Supported – 7210 SAS-X, 7210 SAS-M (network mode), 7210 SAS-T (network mode).

With this release, option is provided to use BFD with static route configured in a VPRN service (IPv4 only). It allows for faster detection of loss of connectivity to the gateway router and disables the use of the route for IPv4 forwarding. If another route to the destination is available, then the best available route (possibly among multiple routes available) will be installed into the FIB and used for IP forwarding.

Refer to the 7210 SAS-M,T,X,Sx,Mxp,R6,R12 Router Configuration User Guide for more information.

9.4.3 Improvement to LAG load-balancing for unicast traffic

Platforms Supported – 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-M (network mode), 7210 SAS-T (network mode).

With this release, option is provided per LAG to select the hashing mechanism used by LAG module for load-balancing flows on the member ports of the LAG. Users can use one of the available options based on the flows they have in their network and select an option that helps improve the load-balancing of flows in their network.

Refer to the 7210 SAS-M,T,X,Sx,Mxp,R6,R12 Interfaces User Guide for more information.

9.4.4 Enhancements to LLDP (Port-ID Subtype)

Platforms Supported – 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE (standalone mode), 7210 SAS-X, 7210 SAS-M (network mode and access-uplink mode), 7210 SAS-T (network mode and access-uplink mode), and 7210 SAS-K.

The following enhancements are made to LLDP in release 9.0R1:

- LLDP now includes the ability to select the port-id-subtype that will be carried in the port-id field. The default remains ifIndex. The ifIndex value is required by some versions of Nokia 5620 SAM to properly build the Layer-2 topology map using LLDP. Changing this value to transmit the ifName or ifAlias in place of the ifIndex may affect SAM's ability to build the Layer-2 topology map using LLDP.
- LLDP System Management Address now includes support for both IPv4 (Address SubType 1) and IPv6 (Address Subtype 2).
- The CLI command "**show system lldp neighbor**" command output has replaced the "**Port ID**" (which printed the ifIndex value) with a new column "**Remote Port**". The "**Remote Port**" column will include the ifDesc (RFC 2863 IF-MIB) when the port description TLV is received. If there is no port-description TLV received or the value is null, the ifIndex will be printed. The "**show port ethernet LLDP nearest-bridge remote-info detail**" command output has been enhanced to print appropriate characters based on the received type.

10 Known Limitations

The following are specific technical limitations that exist in Release 9.0R4 of 7210 SAS OS. The topics are arranged alphabetically.

10.1 ACLs

- At CPM MAF, src-port (front panel port) based filtering does not work for packets received on RVPLS interface. [196845]
- Egress Filter counters does not work for IES service in access-uplink mode SAS-M and SAS-D, however, ACL functionality works fine. [108134]
- The doubly tagged packet ingress on Q.* sap egresses with 3 tags on the Q1.Q2 sap. Such packet not identified as an IP packet as a result will not match the IP ACL on egress. [112850]
- Traffic on dot1q saps can hit egress filter entry attached to “:0” or “0.*” or “*” sap on same port if dot1q traffic matches and matching entry not configured on dot1q sap. The workaround is to configure matching entry with action or an entry with default action forward or drop on dot1q sap. [113270]
- With default action “drop”, if CPU bound packet does not match none of the entry match criteria, these packets are still forwarded to CPU. The workaround for this is, configure a entry criteria with match any and action drop as the last entry in the IP-filter.
- A filter entry action value configured as action should force the filter to pick up the default action configured for the filter. This does not work. It is recommended to explicitly configure the action for each filter entry. [76620]
- MACs that are already learned do not age out after a filter is added to drop packets from those MACs. [73370]
- In an ingress MAC filter policy, the etype and frame-type match criteria do not match packet fields in those packets received with more than one VLAN tag. [72839]
- IP filter match criteria option-present cannot be used to match packet fields for traffic on IES interfaces. [73188]
- Egress SAP statistics and egress ACLs/filters cannot be enabled on a SAP when the port on which it is created has a default SAP configured for 7210 SAS-E.
- In case of 7210 SAS-E, an ingress IP filter policy, Layer 3 and Layer 4 match criteria (such as, src-ip, dst-ip, srcport, dst-port, etc.) do not match packet fields in those packets received with more than one VLAN tag. [72839]

- In case of 7210 SAS-E, a MAC or IP filters applied on a SAP in an IGMP snooping-enabled VPLS service will not be able to block IGMP control packets. [80612]
- From 7210 SAS 6.0 releases, TTL=1 and TTL=255 IP packets would get accounted under SAP stats and these type of packets hit the egress matching ACL's entries.
- When remarking is enabled, remarking occurs on :0,:*,0.* and null SAPs. Accordingly, the egress MAC filter matches that dot1p bit used for re-marking, even though the packets egressing these SAPs do not have a VLAN tag. [161648]
- On SAP ingress use of SAP ingress QoS policy with MAC criteria is mutually exclusive to use of ACLs with IPv6 criteria. [137396]
- Ingress IP ACL cannot be applied for DHCP broadcast packets. [180423]
- The user has to make sure at least 1 ACL entry is free for the ACL re-numbering or copy functionality to work. [164279]
- Time of day policies cannot be combined with IPv6 ACLs.
- In SAS-E Filter logging is not supported. [73135]
- For SAS-X Maximum of 1024 IP Filter and 1024 MAC Filter entries are supported, but CLI allows creation of 8k MAC entries. [105669] [105900]

10.2 CLI

- Non-printable 7-bit ASCII characters (for example, French letters with accents) are not allowed inside the various description fields. These characters were accepted for some description fields prior to Release 5.0. When upgrading to Release 5.0.R1 or later, the user must ensure that the configuration file does not contain any non-printable 7-bit ASCII characters that might have been in any description field prior to Release 5.0. Configurations that do not comply may result in failed config "exec" in CLI and/or during system bootup. User can use "exec -syntax" command to detect if any unprintable characters exist in the current configuration. [99519] [93998]
- The "detail" option for "admin save" command is not supported. Default values under any CLI context can be viewed by using "info detail" command.
- In case of access-uplink 7210 SAS-M, T, D, E, the CLI commands under the context config>router for the management routing instance are not supported. [101636]

10.3 CES

- TDM ports cannot participate in a split horizon group (although it is user configurable). TDM ports is not a supported feature. [101695]
- In Cpipe service of type CESoPSN, if the port is configured for ACR then channel group #1 must be configured in a service and must be operationally up. Channel group #1 is used as the master tributary for deriving adaptive clock using ACR.

10.4 DHCP

- The DHCP broadcast packets are sent to CPU even if DHCP relay is shut down on IES. It is recommended to delete configuration instead of keeping shut down. [161115]
- DHCP packets received over a SDP cannot be identified and option-82 inserted by the node cannot be removed by the node, in the downstream direction. Therefore, if this behavior is not required by the user, the user should not enable DHCP snooping if the DHCP server is reachable over the SDP (either spoke-sdp or meshsdp).

10.5 IGMP Snooping

- In case of 7210 SAS-E, on an IGMP snooping-enabled VPLS service, the 7210 SAS-E does not support multicast forwarding statistics. The show service id service-id mfib statistics command output will always show zero value counters. [81173]
- In case of 7210 SAS-M and T network mode, X, R, Mxp, and Sx in a VPLS service, when IGMP snooping is enabled, the multicast replication is based on Layer 2 MAC addresses.
- In case of 7210 SAS-M and T network mode, X, R, Mxp, and Sx, IGMPv3 is not supported.
- IGMP snooping is not supported for control word enabled SDP in VPLS service.
- In case of 7210 SAS-M and T network mode, X, R, Mxp, and Sx, if single -tagged multicast packets arrive on a null SAP belonging to a VPLS service with IGMP snooping enabled, they are forwarded based on the MFIB if the entry is present, else they are dropped. The same is applicable for double-tagged multicast packets arriving on a Dot1q SAP. [87551]

- IGMP Snooping on sdp with vc type "vc-vlan" requires static configuration.
- On 7210 SAS-M, 7210 SAS-X, 7210 SAS-D, 7210 SAS-T, and 7210 SAS-Sx, default sap (:* sap) does not participate in igmp-snooping process. The following packets are flooded in the service when received on * sap: IGMP General Queries, IGMP Reports/Joins, and Registered and Unregistered Multicast Data Packets. [109045]
- On SAS-E Default sap (:* sap) does not participate in igmp-snooping process. IGMP query sent on default sap is flooded to all the saps and not learned in MFIB. If there is no entry for the group, IGMP report sent on default sap is flooded to all the saps and not learned in MFIB. If there is an entry for the group then report follows the MFIB and data forwarded to only that port.
- Registered Multicast Data packets sent on default sap follows MFIB.
- Un-Registered Multicast Data packets sent on default sap gets flooded in the service. [112586]
- On 7210 SAS-E Flushing out 2047 groups learned in a single VPLS service (by executing a shut on IGMP-snooping command) can cause Dot1ag to flap under loaded CPU conditions. [84904]

10.6 IP

- The 7210 SAS-R BFD sessions created after BFD scaling limit is reached are not coming up even if existing active BFD sessions are deleted, workaround is to shutdown/no shutdown BFD. [181225]
- Unlike other 7210 platforms, the hardware IP route table on 7210 SAS-Mxp, 7210 SAS-R(IMMv2) is shared among IPv4 prefixes, IPv6 /128-bit prefixes and IPv6 /64-bit prefixes. The scaling number of ipv4 and ipv6 /64-bit prefixes varies and more details is available in the 7210 SAS Basic System User Guide under the description for the command max-ipv6-routes.
- "allow-directed-broadcast" is not supported. [122203]
- The ping ip-address detail command should report the interface on which the ping reply was received. This information does not display in the output. [76887]
- IP packets that need fragmentation are not forwarded. However, if the ARP is not resolved for the next-hop, only the first packet is fragmented and sent out as soon as the ARP is resolved. Only CPU-generated packets are fragmented. [76353]

10.7 LAG

- For a LAG configuration with more than one port, every other jumbo frame is dropped. The solution is to increase the MBS from the default value of 128Kbits to 144Kbits for two port LAGs. This is not applicable for 7210 SAS-Sx (standalone mode). [73552]
- For a LAG configuration with more than one port, if a meter configuration does not specify a CBS value, some packets may be marked yellow and be treated accordingly when buffer management begins. The solution is to increase the CBS setting whenever the CIR configured for the meter is greater than 1Gbps. For a LAG with two ports, a CBS value of 64Kbits is recommended. [72497]
- When all the member ports of a LAG are removed and added back, the stats for a SAP on that LAG, belonging to a VPLS or an Epipe service, is reset to zero. [73439]
- In case of 7210 SAS-E, Dot1q tagged LACP packets received on dot1q SAP are dropped instead of forwarding them transparently. [154370]
- LACP and CFM protocol packets count are not shown in the output packets column of the show lag *lag-id* statistics command. Packet count for tunneled LACP pdus are not shown in port statistics. This is not applicable for 7210 SAS-Sx (standalone mode). [77986]
- In an LAG, if port is down due no LACP packets received from other end then unlearned traffic sent on other active ports of lag is also flooded on the oper down port. This is not applicable for 7210 SAS-Sx (standalone mode).

10.8 Management

- The system becomes unresponsive and reboots when the file version check boot.tim command is issued simultaneously from multiple Telnet sessions. Simultaneous execution of this command should be avoided. [76543]
- If an ongoing FTP is aborted, the console and Telnet become unresponsive for a duration that depends on the size of the file being transferred. [76734,74294]
- Max value that can be set for svcVRouterId snmp mib object is 66. [127090]
- SNMP operations on some unsupported SNMP MIBs might succeed.
- snmp dying gasp trap for snmpv1 trap server is not supported. [125543]
- An asterisk "*" indicating an unsaved configuration change may not be displayed after changing some of the parameters under some contexts such as the system and log contexts, (for example, *A:ALU-7210>config>system or A:ALU-7210>config>log). Additionally, the Change Since Last Save field in the show system information command output may not be updated. [61271]

10.9 MPLS

- FRR support with LDPoRSVP or BGP 3107 will not be sub-50ms.
- The operational state of LDP will go down with the reason code "IOM Failure" if the number of /32 prefixes learnt go beyond 1000 for SAS-M, 1500 for 7210 SAS-X, 7210 SAS-R, 7210 SAS-T, and, 7210 SAS-Mxp. [98085]
- A small amount of traffic loss is seen for any MBB event for traffic sent at line rate with a configuration of 200 or more LSPs. [95811]
- Implicit NULL (PHP) must be configured to use LDPoverRSVP tunnel.
- For 3107 L2 services if BGP transport is LDP then LDPoRSVP is not supported, in such configuration sdp will remain down.
- LDP-over-RSVP transport is not supported for BGP SDPs (RFC 3107). SDPs configured in this manner will become operationally up but no traffic will be forwarded. [146172]
- LSR PHP Node copies the tunnel label exp values to single vc-label packets destined to Egress LER but doesn't remark vc-label exp values. [108939]
- FRR failover time for unlearned traffic (such as broadcast, unknown-unicast, and multicast traffic types) will not be under sub-50ms.
- In case of 7210 SAS-R, packet drops are seen during global revert MBB in case of SAP, primary path and FRR path reside on different IMM's. [177662, 177643].

10.10 Mirror

- For 7210 SAS-D, up to 10 Uplink access SAPs can be ingress mirrored. [114247]
- For Access-uplink mode M, T, D, E, when port egress is mirror source, mirrored traffic would contain additional service internal tag.
- For Access-uplink mode M, T, D, E, in case of dot1q mirror destination, when sap ingress is mirror source and single priority tagged packet is mirrored, mirrored packet would not contain priority tag.
- For 7210 SAS-E, mirrored traffic using dot1q sap, profile assignment and hence dot1p remarking does not work. For FC be, profile is ignored and treated as in-profile always. For other FCs, profile is ignored and treated as out-of-profile always. [141252]
- For 7210 SAS-E, packets that get dropped due to egress queue drops are mirrored when port egress mirroring is enabled.

- For 7210 SAS-E, a mirrored packet contains the internal service VLAN ID (that is significant internally to the 7210 SAS-E) when ingress mirroring is configured for an IES SAP. [75852]
- For 7210 SAS-R6, if mirror is enabled on a network port, then mirror traffic carries an extra VLAN tag. This issue exists with SAS-R6 IMMv1 only. [165270]
- In case of 7210 SAS-R6, unlearned service packets egressing out of network port with MPLS header gets mirrored without MPLS header. This issue exists with SAS-R6 IMMv1 only. [174014]
- In case of 7210 SAS-R6, additional mirrored packet sent when both ingress and egress are used as source in case of unlearned traffic. This issue exists with SAS-R6 IMMv1 only. [175208]
- For 7210 SAS-R with IMMv1, egress mirrored frames are copies of the frame as ingressed, any modifications made to a frame at egress, such as VLAN tag, TPID, L3 TTL decrements, and others are not seen at the mirror destination.
- For 7210 SAS-X, only egress rate command can be used on mirror destination. sap-egress qos policy cannot be used with mirror destination (null sap).
- For 7210 SAS-E, packets that get dropped due to egress filters are mirrored when port egress mirroring is enabled.

10.11 OAM

- With EFM OAM, for 100ms timer value it is recommended to use multiplier as 3 (with about 300ms time out).
- Dying gasp SNMP traps takes out-band route when server is reachable through both in-band and out-band, if route-preference configured for in-band which will be ignored.[227153]
- In 7210 SAS-D and 7210 SAS-T access-uplink mode, Down MEP CFM SAA tests Inbound/Outbound values displayed are not proper. [217031-MI]
- On 7210 SAS-D and 7210 SAS-T access-uplink mode, CCM can go in defect state on "send-ais-port-down" enabled MEP when port state toggles. [189432]
- In case of 7210 SAS-R, SAS-T, and SAS-Mxp, if PTP is enabled and the user prefers to go back to NTP time scale, or system free runtime scale, a system reboot is required after disabling PTP, otherwise Y.1731 2DM, 1DM and SLM delay measurements may not be accurate.
- 7210 SAS-M Down MEP CFM packets does not follow QoS policies.
- In case of 7210 SAS-R, Up MEP on SAP in Epipe Service CFM reply messages egressing from SDP binding are sent with exp value 7 when remarking is disabled. [202804]

- CFM reply packets are always taken as in-profile while re-marking even when request packet is classified to out-profile. [202687]
- Up Mep on Null or * sap with egress sap dot1q/qinq, following is the behavior. [202717]
 - CCM message: by default Dot1p bits set to 0. If remarking is enabled in egress sap then takes configured dot1p.
 - Other messages: EPIPE service in 7210 SAS-R6, 7210 SAS-X, 7210 SAS-M, 7210 SAS-Sx, and 7210 SAS-T (network mode only) by default Dot1p bits is set to 0. If remarking is enabled in egress sap, then the configured dot1p is taken.
- In Testhead jitter and latency values are provided only when the test traffic encapsulation is same as test SAP, for example, Q1.Q2 doubly tagged traffic is used if test sap is of Q1.Q1 encapsulation. [152369]
- On 7210 SAS-D, Mirror and Testhead functionality cannot be configured at the same time. [163084]
- For 7210 SAS-D, the show eth-cfm mep mep-id domain md-id association ma-id command will not display CCM ERROR, CCM XCON frames in its output.
- For 7210 SAS-D, the show eth-cfm mep mep-id domain md-id association ma-id remote-mep rmep-id command will not display some TLVs and details.
- On the 7210 SAS-D platform, even if a SAP is administratively shutdown, the hardware state machine receives and processes the CCM packets sent by remote peers and the CFM MEP remains up and shows no CCM defect.[126719]
- Under a high CPU load, ETH CFM defect reporting and clearing may be delayed.
- All platforms except SAS-R6, when a SAP has UP MEP configured on it, the CFM frames would use the forwarding path of the service. This causes the "Ingress Stats" of SAP Statistics to increment when UP MEP sends packets. These packets are also counted in "Ingress Drop Stats", if SAP's "statistics ingress counter-mode" set to default that is, "in-out-profile-count". [141370]
- In case of 7210 SAS-D, M, and 7210 SAS-T down MEP, the 7210 SAS does not display the remote MEP's MAC address in the display output of the CLI show command "**show eth-cfm mep mep-id domain md-id association (ma-id) all-remote-mepids**".
- In case of 7210 SAS-D, M, T, the ETH-CFM CCM Sequence Check for out-of-sequence CCMs and Last CCM Sequence Number tracking is not done for received CCMs.
- In case of 7210 SAS-D, M, T, the ETH-CFM CCMs with incorrect DMAC are not dropped by hardware.
- In case of 7210 SAS-D, M, T, the ETH-CFM CCMs transmitted will always have the Sequence Number as zero.

- In case of 7210 SAS-D, M, T, the ETH-CFM CCMs with invalid Data TLV-s, Invalid Port/Interface TLV-s do not cause CCM Errors
- In case of 7210 SAS-D, M, T, the show eth-cfm mep *mep-id* domain *md-id* association *ma-id* command would not display CCM Tx Count.
- Before configuring any port as "loopback-svc-port", it is always recommended to remove any configurations made on this port, and ensure it has only default configurations.
- Test head test port is not supported when port mode is configured as access uplink (these ports show up in show port display as "I2up"). [150123]
- Test head marker packets etype is always set to 0x0800, user defined etype in test head profile is not carried in marker packets. [153288]
- Test head marker packets does not carry any Layer 4 info defined in test head profile, any qos/filter applied on a sap based on Layer 4 fields will not take effect for markerpackets. [153414]
- For modifying saps from service, internal mac swap loop back port, test head loop back port it is recommended to deconfigure mac-swap loop back configuration. [150496, 150403]
- With UP MEP, removing all lag member ports "defMacStatus" is not reported, "defRemoteCCM's" on both ends are reported. [136119]
- When SAP in the parent MVPLS instance goes into STP blocked state, the associated SAP in the child VPLS instance is shut down. This prevents any packets from being generated in the host path. In 7210, CCM-s are hardware generated and hence the CCM state m/c works properly. However, the LB-s and LT-s, which are generated and processed in software, stops working. [109722]
- In a spoke SDP with the control word configuration enabled, vccv-ping from the remote end does not return a response when the LSP is shutdown. [80905]
- If SAP is configured as static mrouter port, port loop back with mac swap does not work in unregistered multicast traffic. Unregistered multicast traffic that is sent out of the SAP is not looped back but registered/learned multicast is looped back. [130327]

10.12 Routing

- "configure router bgp family" by default is set to "ipv4". It needs to be configured to "vpnipv4" for VPRN configuration.[122553]
- When the BGP routes fails to get installed due FIB is full, BGP Peer goes operationally down.
- OSPFv2/v3 would operationally go down on exceeding the FIB limit.[144137]

10.13 QoS

- The P2MP LSP traffic received on network interface uses only unicast meters though multicast meters are associated with that FC.[230824]
- VPLS/EPIPE for SAP to SAP packets on hybrid port DSCP remarking is not supported. [225903]
- In case of 7210 SAS-D, DSCP can be remarked only for packets received with number of tags matches to tags configured on ingress SAP.
- Access-Egress policy enabled with dscp remarking will remark DSCP bits in data packets egressing out of L2 SAPs (VPLS/Epipe/I-BVPLS/B-VPLS). It is recommended not to enable DSCP remark for ports carrying L2 SAPs. [121134]
- In case of 7210 SAS-R, when SAP remarking is not enabled for L2 service SAPs and port remark is enabled, Port Remarking for L2 service does not work if ingress is null/*/0/cp-1 SAP and egress is dot1q/q1.*/q1.q2. or Ingress is spoke-ether and egress is q1/q1.*/q1.q2 SAPs or svc-sap-type qinq-inner-tag-preserve. Workaround is to use SAP egress policy for remarking. [170807,168509]
- In case of oversubscription of 7210 SAS-R IMM bandwidth, higher priority packets may get dropped.
- For 7210 SAS-D from 6.0R2 release, 7210 SAS-D per port total available buffer is 146KB compared to 157KB in previous releases. The shared portion of buffer available has reduced from 89KB to 78KB. [159987].
- For 7210 SAS-D, with a slope policy, queue depths for lower priority queues do not get limited to max average when the scheduler is in strict mode and if lower priority queue is not serviced by scheduler due to high traffic on higher priority queue. [111651]
- Meter buckets are re-initialized when the rate value is modified. Rate values are modified by explicitly changing the rate values using the appropriate CLI command or by changing the adaptation rule. [84395]
- When a port is congested, a small amount of excess traffic is sent out of the lower priority queues. The amount of excess traffic depends on the packet sizes. [111644, 111664]
- For 7210 SAS-E, there is a small difference in the actual rate of traffic egressing out of a port for a given egress rate limit value. The difference in the actual rate is affected by the size of the packet. For example, it is noted that for an egress rate limit value of 40Mbps, the difference is approximately 2Kbps to 150Kbps for packets sizes between 100 bytes to 9212 bytes.
- A maximum of 15 network QoS policies of type ipinterface with unique mapping of FC to EXP values can be created, and these policies must be shared among 32 IP interfaces.

- If a network QoS policy with classification based on match of a Dot1p value '0' is associated to a network port, which has IP interfaces using either a null or a dot1q:0 encapsulation, any untagged IP packet received on the network IP interface will get classified to the FC designated by this rule. The same behavior is applicable to a null SAP when it receives untagged packet.[98819]
- With a slope policy, queue depths for queues 1 to 7 for a port do not get limited to max-average when the scheduler is in strict mode. [85063]
- If an IP interface is configured on a hybrid port, following is the behavior.
 - IP interface on q.* encap, single tagged packets gets proper classification and policing
 - IP interface on 0.* encap, up to two tag packets gets classification and policing properly as per the config
 - IP interface on q1.*, and q1 interface, q1.q1 packets gets proper classification, but q1.q2 does not
- PBB: Ingress classification on B-SAP using I-TAG PCP bits and egress remarking of ITAG PCP bits is not supported. [123642]
- In case of 7210 SAS-X and 7210 SAS-R, for VPRN L3 SAP, per SAP egress remarking is not supported. It is recommended to use port based egress remarking. If DSCP is used with port based egress marking, then L2 SAP traffic is also marked. Therefore, when having a mix of L2 and L3 SAPs on the port, it is recommended to use only Dot1p based marking.
- In 7210 SAS-R CIR oversubscription is not supported both at Aggregate Level and Port Level (ERL). Queue and Aggregate CIR rates should be configured such that bandwidth is available to each level to service all Committed Information rates at each level. [175278]
- During congestion in the system , and if a queue is not able to get shared buffers to hold at least number of packets equal to weight configured, then the configured weights cannot be guaranteed. [163712]
- For 7210 SAS-X, the overridden CBS of a queue can only be greater than the CBS defined in the sap-ingress qos policy. [160811]
- For 7210 SAS-X, when an ERL lesser than 700Mb is applied on a 1G port and when two queues are in the PIR scheduling loop, traffic is seen in low priority queues even though there are drops in high priority queues. [152374]
- When remarking is enabled on access egress for a Dot1q port, the Dot1p bits in the outer customer tag get remarked when the traffic is sent out of a Dot1q default SAP. Similarly, when remarking is enabled on access egress of NULL encap port, Dot1q bits in the outer customer tag get remarked when traffic is sent out of a NULL SAP. The workaround: In the case of a NULL SAP, remarking can be disabled to preserve the Dot1p bits.[86818-MI].

- All untagged packet received on null access sap or null encap network port is classified to a FC that is associated with an entry with a match criteria set to dot1p value of "0". [98819]
- For 7210 SAS-X, Queue CIR/PIR minimum value can be set to 26kbps. Configuring lesser value results in traffic stops on egress.[102578]
- In case of 7210 SAS-X, higher Cir-Level queue's PIR traffic may affect lower Cir-level queue's CIR traffic when port egress-rate or sap egress agg-rate-limit configured. Similar issue of higher CIR queue rate affected by lower CIR queue rate seen, if queue's at same CIR level and difference between CIR values of the queues is larger. This issue is seen where egress-rate or sap egress agg-rate-limit configured value closer to sum of queue CIR's value configured. [111002, 111657,126359]
- In case of 7210 SAS-X, egress queue drop counters count only tail drop packets and WRED drops are not counted. [109298]
- For 7210 SAS-X, when 2 or more SAP queues are in the same PIR level but having different weights, traffic distribution may not be as per PIR weight of the queues. This behavior was observed in case of higher bandwidth allocated to one or more higher level queues of the same SAP.[107013]
- For 7210 SAS-X, when Traffic is egressing out of network port for VPLS/Epipe services, traffic is rate-shaped without taking MPLS or dot1q header into account on network-port. This leads to actual traffic throughput more than configured queue rate (CIR and PIR). The workaround is to take the network-header into account while configuring the queue CIR and PIR. This behavior is also with egress-rate configured on the port. In case of vpn service, MPLS header accounted for rate-shaping or egress-rate.[101713]
- For 7210 SAS-X, when traffic is sent from dot1q-sap to null-sap, the actual throughput is less than the configured queue CIR/PIR. The workaround is to take the dot1q-header into account while configuring the queue CIR and PIR. This behavior is same with egress-rate configured on the port. [102176]
- Only the outermost tag is marked with dot1p bits when remarking on egress is enabled for QinQ access SAPs. [108010]
- When configuring network QoS policy using SNMP, it is recommended to first create policy specifying type of policy (ip-interface or port type) and then modify policy to set user defined or default values of the policy. [123171]

10.14 Security

- 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.

- 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 7210 node 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.
- If the TACACS+ start-stop parameter option is enabled for accounting, every command will result with 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 has the 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 define if the user has only console or FTP access. The 7210 SAS OS supports a default template for all TACACS+ authenticated users.
- If TACACS+ is first in the authentication order and the TACACS+ server is not reachable, the 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 box. Note that inconsistencies can arise depending on certain combinations of local, RADIUS and TACACS+ configurations. For example, if the local profile restricts the user to only FTP access, the authentication order is TACACS+ before local, 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]
- If a source-address entry is configured for inband RADIUS servers, the source address (IP address) is used as the NAS IP address, otherwise the IP address of the system interface is used.
- In defining RADIUS Vendor Specific Attributes (VSAs), the TiMetra-Default-Action parameter is required even if the TiMetra-Cmd VSA is not used [13449]

10.15 Timing

- On 7210 (except SAS-K platform), NTP time correction can't happen if system time is set to beyond +/- 32 years of NTP time.
- In case of 7210 SAS-D fixed copper port as SyncE reference is supported with port speed as 1 Gbps and Auto negotiation is ON or Limited.

- Ethernet ports that use dual-rate fiber SFPs cannot participate in a Synchronous Ethernet Network.
- Applying the command "debug sync-if-timing" on a second qualified reference has no effect.
- SyncE Reference switch is based on LOS and not based on signal degradation.
- Ethernet ports that use copper SFPs cannot participate in a Synchronous Ethernet Network.
- Standby CPM takes around 10-15 seconds to update sync-if-timing status after CPM switch over. [169225]

10.16 Services

- In case of 7210 SAS-Sx, when null/:0/*/:0.* is configured on port, control protocol source mac, example LACP packet mac is learnt in FDB table.
- FRR and Eth-ring switchover timings may not be within 50ms when copper ports are in use.
- It is NOT recommended to configure MC-LAG without LACP and MC-LAG with LACP on same node, it may take more for traffic to converge in some failure cases. [192846]
- In case of PW switching, mismatch in control-word configured on static segment results in packet corruption. [186952]
- Multicast data packets with ttl=1 are not forwarded in service when MVR is configured. [143559]
- SHG can be configured either in "mvr vpls-service" or "user vpls-service", but not in both.
- The system accepts packets with sizes exceeding the port MTU by 4 bytes, as listed in below scenario if egress port has proper MTU to transmit these extra bytes.
 - 1522 byte (includes 4 byte FCS) single vlan tagged or 1526 byte (includes 4 byte FCS) double vlan tagged packets received on null sap is forwarded where ingress is null port and egress is dot1q/qinq port.
 - 1526 byte (includes 4 byte FCS) double vlan tagged packet received on dot1q sap is forwarded where ingress port is dot1q and egress is qinq port. [75221]
- A MAC will not age out as long as STP BPDUs from that source are received, although data traffic is not present. [71658]

-
- When qinq etype (x) is configured on a port and a sap q1.q2 is created, q1.q2 tagged traffic mapped to this sap only if the outer tpid is (x) and inner tpid is 0x8100. If q1.* sap exists on same port then traffic mapped to q1.* sap if the outer tpid is (x).
 - All R-VPLS interfaces by default chassis mac is assigned, different mac can be configured using CLI "interface <intf-name> mac". [139425]
 - R-VPLS service not supported with *.* Saps. [140881]
 - For R-VPLS service, mac address resolved by ARP is also learnt in L2 FDB table; hence ARP ages out upon FDB entry age out.
 - In R-VPLS operational IP MTU is set to least of participating SAP's MTU. SAP MTU is Port MTU – DLC header, wherein DLC header is 14 Bytes for NULL port, 18 Bytes for dot1q port and 22 Bytes for qinq port.[141391]
 - A service configured with 'svc-type any' and null-sap to null-sap allows packets up to 1514 including 4 byte FCS if the port MTU is default value of 1514. Workaround is to increase the port MTU from default value of 1514. [116239]
 - In VPLS and VLL services, port MTU checks are performed only at the ingress. [92910]
 - In case of access-uplink mode SAS-M, T, D, E, an Epipe service, traffic is not switched if the source MAC is a multicast or broadcast address [71437]
 - For 7210 SAS-E, when a ARP request is received for an address configured for an IP interface in a VPLS service, the first ARP reply is sent to all the SAPs, instead of the SAP on which the request arrived. [94288]
 - For 7210 SAS-E, MAC address learning rate is slow for certain sized packets. The average learning rate is approximately 200 MAC addresses per second. The learning rate for packets with sizes of 80 and 260 have been determined to be lower than this average. [77067]
 - For 7210 SAS-E, Layer 4 load balancing for broadcast, multicast and unlearned unicast traffic is not supported. [72425]
 - In 7210 SAS-M configured in network mode, temporary data loop around 10-20 msec during revertive mode of 8032. [129307]
 - Source BMAC, B-SA learnt in B-VPLS service starts aging only after C-SA associated with B-SA are aged out in case of I-VPLS. In case of PBB-EPIPE, B-SA is not aged out once it is learnt.
 - PBB I-tag etype is not configurable, Its value is 0x88e7. PBB B-Tag etype is not configurable Its value is 0x8100
 - Traffic from a B-VPLS SAP to a I-VPLS SAP matches the p-bit on the B-header instead of the customer dot1p bit when a MAC filter is applied on the B-SAP. [158408]

- For network mode 7210 SAS-M, 7210 SAS-T, 7210 SAS-X, 7210 Mxp and 7210 SAS-R6, processing of ingress BUM traffic on a SAP or spoke-SDP configured in a VPLS service uses up port egress resources. This results in traffic drops for egress traffic out of the port on which SAPs or spoke-SDP is configured, if BUM traffic contributes to greater than or equal to 50% of port bandwidth. Depending on the packet size and ingress BUM traffic rate, different amount of drops are observed on different platforms. This issue is not applicable to 7210 SAS devices configured in access-uplink mode. [85380, 215022]
- PBB packets with UCA (Use Customer Address) bit set are not forwarded. [113909]
- For epipe service of "svc-sap-type qinq-inner-tag-preserve", the priority of the first tag in the packet egressing will always be 0 unless remarking is enabled.
- For epipe service of "svc-sap-type qinq-inner-tag-preserve", Ingress and Egress IP classification and filter will not work if more than 2 tags are received.
- Traffic that hits a blackhole route in a VPRN Service, would not get accounted under filter statistics. [125527]
- Fragmentation is not supported for IP Packets received on L3 SAP of vprn service [121145]
- When packet is received on a SAP, the service MTU check includes the length of the packet and the SAP delineation encapsulation overhead (that is, 4 bytes for a dot1q tag or 8 bytes for a QinQ SAP). Similarly, when a packet is received on a SDP Binding (also known as PW) is of type vc-vlan, the service MTU check includes the length of the encapsulated packet along with the vc-vlan encapsulation length. If the packet length is greater than the service-mtu, that packet is dropped.
Note: This limitation is not applicable to 7210 SAS-K. On 7210 SAS-K, service-mtu includes the length of the packet excluding the service delineation. For more information, see the Services User Guide.
- With service-mtu-check disabled, a null SAP allows 4 bytes more than the port MTU for tagged traffic (excluding FCS); a dot1q SAP allows 8 bytes more than the port MTU for tagged traffic (excluding FCS); and a dot1q-star SAP allows 4 bytes more than the port MTU for un-tagged traffic and 8 bytes more than the port MTU for tagged traffic (excluding FCS).
- In case of 7210 SAS-X, max 63 saps can be configured on Hybrid port. [150473]
- SAS-X only - Discard unknown on b-vpls service is not supported. [122038]

10.17 Statistics

- SAP ingress meters (counters) are incremented for packets dropped by a filter on that SAP. [70878]

-
- Packets with CRC errors are accounted for in the ingress meter calculations. [80966]
 - Packets discarded as the result of a discard unknown-source and discard-unknown configuration are accounted for in the ingress meter calculations. [84842]
 - Egress sap statistics not supported for VLAN range SAPS
 - On access ports, protocol packets of EFM, LACP, Dot1x, and DWL are counted as part of SAP statistics if NULL SAP, Dot1q Explicit NULL SAP, or Dot1q Default SAP is configured on an access port. [95361]
 - Non-routable traffic received on access IP interface not accounted in sap stats. [137643]
 - For 7210 SAS-E, for SAPs, if egress filter and egress statistics are enabled together then egress filter counters cannot be used to obtain count of packets matching the egress filter entry. However, egress filter functions appropriately. If the user needs to obtain a count of packets matching the egress filter entry, the egress statistics must be disabled with the no packetsforwarded-count command in the SAP context where the egress filters are in use. [93524]
 - For 7210 SAS-E, Accounting records can only count packets or octets at a given time. This is configurable by the user. The configuration is also used to change the behavior for statistics collection.
 - For 7210 SAS-E, the monitor service id service-id sap sap-id rate command displays statistics in either packet or octet mode based on the accounting record configured. The utilization rate shows the appropriate values only when operating in octet mode
 - TTL=1 and ARP packets are not accounted in a vprn L3 SAP Statistics.
 - Packets larger than port MTU are learned and are accounted for by rate limiters. However, they are dropped as expected. [73497]
 - Accounting statistics for a spoke SDP in a VPLS service show extra egress packets when the destination to which the traffic stream is being sent is already learned, but actual packets on the wire are correct. The percentage of error depends upon the packet sizes, FDB size, rate of traffic and the duration of the traffic.[81608/94306]
 - IGMP packets sent out of a SAP or spoke SDP are not accounted in SAP or spoke SDP egress statistics. If they are received from a peer and forwarded out of a SAP or spoke SDP, they are accounted. [88332]
 - Ingress SDP statistics are accounted against the primary spoke SDP, even if traffic is received on other secondary spoke SDP. [93627].
 - When continuous traffic is flowing through sap, "Ingress Drop Stats" may not display proper results if sap "statistics ingress counter-mode" is "in-out-profile-count"

- In 7210 SAS-R, Queue drop statistics on dot1q/qinq L3 VPN sap case, vlan tag length of 4byte is not considered for octet conversion.
- In case of access-uplink mode 7210 SAS-M, SAS-T, SAS-D, egress ACL stats are not working in a null-star service for null Star,0.* sap,:* sap,:0 sap, however ACL functionality works. [122804]
- For 7210 SAS-X, CPU traffic bypasses egress scheduler policies and is not counted in egress queue counters statistics. [105180]
- For 7210 SAS-X, queue rate calculation is frame based. When traffic egress or ingress out of a queue, the configured queue rate accounts for IFG and preamble of the frame. Port egress rate calculation as well is frame based.[103790]
- For 7210 SAS-X, egress queue stats "Octet Count" output may not count network-header bytes in case of network port or dot1q header bytes and in case of null to dot1q sap. [109272]
- In case of 7210 SAS-X BUM traffic received on a sap are accounted in egress queue counts of the source sap. [106158]

10.18 STP

- STP BPDUs received on uplink sap hits the unicast meter instead of multi-point meter. [128453]

10.19 System

- The below mentioned issue and workaround is applicable to SAS-M and SAS-E units manufactured before 15-Dec-2013. For any alarm-box with open-circuit-voltage greater than 12VDC, the ESD/lightning protection circuit inside SAS-M or SAS-E may cause unexpected leakage current on the alarm contacts output. To overcome this design limitation, the customer can add ¼ watt resistors (matching resistor values on the three wires) to limit the leakage current. Depending on the remote device, try these resistor values from high to low in the following sequence: 33K, 22K, 16K, 9.1K, 6.8K, 5.1K and 3.9K.
- It is recommended to shutdown 10 Gig MDA before physically removing it from 7210 SAS-M chassis. If 10 Gig MDA is removed without shutdown in some cases fixed 10 Gig ports 1/1/25 and 1/1/26 flap. [144816]For SAS-D and SAS-T, the Fixed Copper ports cannot operate in 1Gig mode with auto negotiation disabled. [75345, 113818]

-
- When performed "file dir" on empty drive (cf1 or uf1), shows "MINOR: CLI File Not Found "uf1:." [178197]
 - After reboot for about 15 minutes, card temperature readings are not displayed in 7210 SAS-D.
 - OAM DNS lookups do not work correctly unless the full DNS name is provided. [54239, 54689]
 - Inserting and removing SFPs in rapid succession causes the "SFP/XFP Checksums do not match" message to be displayed on the CLI session. If this message appears, execute a shutdown command followed by a no shutdown command on the offending port to resolve the issue. [76935]
 - BOF password configured by user gets reset in case 7210 rebooted with "boot.tim" version (prior to 4.0) which does not support BOF password. [145037]
 - The number of files in the root directory is limited to 100. As a possible workaround, create a directory in the root directory and use that to save/store files. [75227]
 - There is currently no show command to display the current values of password hash settings. [32747]
 - When the password aging option is enabled, the reference time is the time of the last boot, and not the current time. The password expiry is also reset on every reboot. [64581]
 - A port LED may glow if a 1Gig fiber SFP is inserted (without connecting a cable) with 100Mbps speed configured. (Note: 100Mbps mode is not supported for 1Gig Fiber SFPs excepting dual-rate fiber SFPs and copper SFPs). [85620]
 - After 497 days, system up-time will wrap around due to the standard RFC 1213 MIB-II 32- bit limit.
 - Dual rate SFPs (3HE04116AA and 3HE04117AA) connected to GigE SFP require autonegotiation to be enabled to operate in 1G mode. [78737]
 - **default.cfg** is a file name reserved by the system. Do not create a file with this name in the root directory. [76972]
 - CLI "file dir" does not start to list when number files are more than 700+ files on flash cf1
 - If an ongoing FTP is terminated, the console and Telnet are unresponsive for a duration depending on the size of the file being transferred. [76734, 74294]
 - For 7210 SAS-E, the link state of fixed copper ports remain in their current state even after a system reboot is initiated, but traffic is not forwarded. They are initialized during system startup. [76466]

- The user is recommended to protect the out-of-band (OOB) Ethernet management interface to avoid high CPU utilization when high-rate of traffic (For example, High rate of traffic due to DoS attach, high rate of broadcast traffic due to network miss configuration, and so on) is received on that interface. The system does not rate-limit ingress traffic on the OOB port and users should use other mechanisms to achieve this process. Management Access filters can be used to filter traffic destined to CPU. It uses CPU resources and cannot be used to drop high-rate of ingress traffic on the OOB port.
- It is recommended to shutdown external CF or USB or SDCARD using CLI "file shutdown <>" before they are physically removed from the system.
- In case of 7210 SAS-M, ports used as "no-service-ports" in bof should not be used in configuration, else execution of configuration file errors out. [120103]
- LED's on vwm-shlef controller turns "Green" when configured shelf-id matches with rotary-id, but sometimes it shows "Amber". Executing "show system vwm-shelf" OR "show system vwm-shelf id" turns LED to "Green" if configured shelf-id matches with rotary-id.
- Quick insert and removal of USB may report flash-device failure, allow 15-30sec settling time to avoid these flash-device-false failure messages.
- When the show service fdb-mac command is executed through the console while the MDA is initializing (and when traffic is coming into the box), a software crash could occur. To avoid this anomaly, wait until the MDA is initialized and the ports are up before issuing the command. [74051]

11 Known Issues

The following are specific technical issues and limitations that exist in Release 9.0R4 of 7210 SAS devices. The topics are arranged alphabetically.

NOTE:

- 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.

11.1 ACLs

- Egress filter counter does not increment for action drop on 7210 SAS-R IMMv2 and SAS-Mxp, however filter functionality works fine. [190140-MI]
- IPv4 or IPv6 ingress ACL configured on network interface filters matching MPLS encapsulated IPv4 or IPv6 service traffic. [199826-MI]
- For 7210 SAS-E, egress Packets not be dropped when IP filter is configured on QinQ Uplink SAP with QinQ etype is configured to non default values on the port. [112881]
- In 7210 SAS-E, egress Ip/Mac filters action not work on 0.* SAP for Epipe and Vpls service. [114263]
- For 7210 SAS-E, ingress and egress filters do not block STP packets when STP is enabled in a service. [75921]
- For 7210 SAS-E, a **mgmt-access-filter** with a **deny-host-unreachable** action sends "Destination Net Unreachable" instead. [73676]
- For 7210 SAS-E, an egress filter applied to drop all traffic on a port also drops EFM loopback traffic. [82782]
- For 7210 SAS-E, an IP filter applied at SAP ingress to filter IGMP packets does not work.

11.2 CLI

- In case of 7210 SAS-R, CLI rollback for PTP configuration is not supported. [219874-MI]
- Following are some of the known issues with CLI Rollback on 7210 SAS-R:
 - ERROR seen while doing Rollback revert for a management access filter. [188689]
 - Rollback failed for maximum routes in VPRN. [187417]
- There are some unsupported CLI and CLI options displayed in the CLI command set.
- Some of the show, monitor, and tools CLI command output displays unsupported fields and modules.
- The output of the **tools dump** command is not aligned properly when issued from a Telnet session. [76876-MI]
- The 7210 SAS-R6 does not support BGP RR functionality, though the CLI is available.
- 7210 SAS-X, CLI "monitor port" output is not 100% accurate all time. [110978-MI]

11.3 CES

- Commands executed quickly after provisioning a T1/E1 MDA would get delayed response due to the MDA initialization (given that the commands needed IOM/MDA routines to be called such as **show port**). [94277-MI]
- SAP ingress statistics in a CPIPE service does not increment when LOS is reported on a DS1/E1 port. [98828-MI]
- Provisioning m24-100fx-1gb-sfp MDA (which is pre-provisioned as MDA1) on MDA2 can lead to undesired behavior especially when a CES card is installed in the MDA 2 slot. [97786-MI]
- TDM ports cannot participate in a split horizon group (although it is user configurable). It is not a supported feature. [101695-MI]

11.4 IGMP Snooping

- On 7210 SAS-K, null tagged IGMP snooping control packets are dropped on a `:* sap` [226117 – MA]

- In case of SAS-R, in a VPLS service with IGMP-snooping enabled and mrouter configuration on multiple endpoints, very few packet drops may be seen on the mrouter endpoints, whenever a new multicast group is created or deleted due to a first member addition or a last member leave.[201239-MI]
- In an IGMP snooping enabled VPLS service, if the IP interface is removed and added back, multicast traffic between spoke SDPs may not resume towards one or more spokes. The spoke SDPs mentioned here are configured for the M-router port. The work-around is to remove the static M-router configuration and add it back; or execute a shutdown command and a no shutdown command of the spoke SDPs. [91856-MI]
- In a scaled setup with spoke SDPs configured as M-router port multicast traffic for some services is not forwarded over some spokes when the LAG flaps. Remove and reconfigure the M-router port for the spoke SDP restores the traffic. [97851-MI]
- In 7210 SAS-E, when more than 2047 (S,G) joins are received in a scaled setup, it might result in hash collisions in the multicast forwarding. In such cases, learned groups can be removed/added back, with messages appearing on the console. (95345)
- On an IGMP snooping-enabled VPLS service, the 7210 SAS-E does not support multicast forwarding statistics. The **show service id service-id mfib** statistics command output will always show zero value counters. [81173]

11.5 IP

- In case of 7210 SAS-R, for IPv4 or IPv6 route learnt through BGP, sometimes next hop shown as junk value in “show router fib” output. There is no issue with traffic forwarding.[213748-MI]
- In case of 7210 SAS-K, if IP packet forwarded through R-VPLS requires fragmentation then such packet will be not forwarded, also ICMP error message also not generated.[195127-MI]
- For RVPLS IP forwarding ARP is closely bound to FDB entry, ARP entry will be removed if MAC in FDB ages out. Traffic will not be forwarded till ARP resolved. To avoid traffic loss it is recommended to configure FDB timeout be greater than or equal to ARP timeout. This issue is not applicable for 7210 SAS-K. [190982-MI]
- SNMP query of the vRtrActiveArpEntries object does not return the correct value. The CLI reports the correct number of ARP entries. [80788-MI]
- 7210 SAS-E does not support secondary IP interfaces. However, these are configurable through SNMP, and should be avoided. [76848].

- When ARP is cleared, few packets from a single flow IP traffic gets load balanced in case ECMP is enabled. [156758-MI]
- With ECMP enabled, mac-ping and eth-cfm loopback test fails for LDP based SDP bindings. [161380-MI]
- 7210 SAS does not support Sub Second hello timer for VRRP.
- In 7210 SAS-T and 7210 SAS-R6 with IMMv1, if IPv4/UDP or IPv6/UDP packets with UDP destination port value equal to 3784 and IP TTL value NOT equal to 255 or 1, that are received on any L2 service or L3 interface are not forwarded. [184471-MA]
- In 7210 SAS-R, directly connected active routes take considerable amount of time to come up. [171954-MI]
- In 7210 SAS-R, reducing the interval/timeout timers much below default values is not recommended for OSPF, IS-IS, BGP, LDP and RSVP to ensure stability under transitional events like a CFM switchover. [56792, 58891-MI]

11.6 LAG

- For 7210 SAS-K, it is recommended to have max of 2 LAG and 2 ports in Uplink LAG even though CLI allows to configure. [207440-MI, 207653-MI]
- With LSP over lag, traffic is sent out on a port added to the LAG sub-group which is in stand-by mode, workaround is to shut/no shut of LAG. [159334-MI]
- If in a two-port LAG containing two sub-groups with one port each, a **shutdown** on the port belonging to the active sub-group will flap the LAG. [85967-MI]

11.7 Management

- SNMP walk of vRtrConfEntry shows the VPLS-management and management instance as active even though these are not currently supported. [76832-MI]
- SNMP query of the following operational rates does not return the correct values. The value returned is 0. CLI reports correct operational values. [76853-MI].
 - tAccessEgressQueueOperPIR
 - tAccessEgressQueueOperCIR.
- If a source-address is configured for NTP, and if the system is rebooted with an older time set (using **admin set-time**), NTP takes a few iterations to synchronize for the first time. [86897-MI]

- The CLI allows the user to specify a TFTP location for the destination for the **admin save** and **admin debug-save** commands which will overwrite any existing file of the specified name. [18554-MI]
- The 7210 SAS does not support storing more than 500 events for log destinations memory and SNMP. Although, the CLI and MIB allows up to 1024 to be configured, it is recommended not to exceed 500.
- SNTP broadcast packets are not processed when they are received with an all-ones address. They are processed if they are received with sub-net broadcast address. [73662]
- LACP flaps when starting an SSH server (no shutdown of SSH server). [75648]
- If there are NTP broadcast client configurations over the **management** routing-instance, and if the out of band Eth-management port is disabled, on configuring an NTP server and removing it will remove the broadcast client configurations as well. [101255-MI]

11.8 MPLS

- For 7210-R, packet drops are observed in VPLS/EPIPE service when traffic switched over to MBB secondary to primary path. [234414-MI]
- MPLS packets received with multicast destination MAC with ingress label belongs to a P2MP tunnel are not processed. [226290-MI]
- Packet drops are observed in VPRN service when traffic switched over to MBB path. [234321-MA]
- In case of 7210 SAS-R and 7210 SAS-Mxp, FRR timings are greater than 50ms when saps are also configured on a hybrid port. [178443-MA].
- In case of 7210 SAS-R, modifying the LDP hello timers while the hello adjacency is up does not come into effect, until the adjacency bounces. However, after two High Availability switchovers, the active CPM or CFM starts using the new timer value. [112617-MI]
- Tools perform router MPLS CSPF command does not accept SRLG group name having alpha characters [155853-MI]
- Tools perform router MPLS command does not work for SRLG group names having numeric character length of more than 10. [156220-MI]
- The **show router rsvp interface *interface-name* detail** command displays incorrect Auth Rx Seq Num and Auth Tx Seq Num values. [86903-MI]

11.9 Mirror

- In case of 7210 SAS-K 2F4T6C, Mirroring of single labeled packets are malformed. [246688-MI]
- 7210 SAS-Sx (standalone mode) remote mirroring does not work for EFM, LACP and LLDP packets. [228812-MI]
- 7210 SAS-Sx (standalone mode) Dot1q/QinQ mirror traffic dot1p bits are not getting remarked according to the remark policy applied. [228752-MI]
- 7210 SAS-Sx (standalone mode) mirror egress stats does not reset to 0 after removing and restoring the sap when there is no source configured.[227838-MI]
- 7210 SAS-R6, egress mirror also mirrors ingress ip multicast data packets. This issue exists with SAS-R6 IMMv1 only. [189812-MI]
- Ingress QoS policies applied for forwarded traffic will also be reflected on its mirrored traffic, if mirror destination is a null SAP. [73951-MI]
- Log events are not generated for mirror application. [72100]

11.10 OAM

- * In case of 7210 SAS-Sx, Y.1731 2-DM values are not proper if "ntp server ptp prefer" configuration is used. [251809-MI]
- * In case of 7210 SAS-K-12, EFM and LACP packets are getting dropped when filter is attached to ingress sap (null & * saps) on that port with either "default-action drop" or action drop of any matching entry. [249026-MA]
- * In case of 7210 SAS-Sx/S, Port loopback with mac swap does not work if 10G port is used for mac-swap. [248424-MI]
- In case of 7210 SAS-K 2F4T6C,CCM not converged on standby spoke sdp MEP configured on VPLS service. [247446-MI]
- In case of 7210 SAS-K 2F4T6C, CCM not converged when force-vc-vlan-forwarding enabled with non-default dot1q-ethertype.[247340-MI]
- CFM OAM tests not works for all MEPs of multiple MEPs of different SAPs, where SAPs belong to single VPLS and SAPs are configured on single port. [236788-MI]
- 7210 SAS-K, when multiple switchover tests are carried out, it is observed that few times Eth-Ring switch over times are greater than 50ms. [221913-MI]
- 7210 SAS-K, packets generated by testhead OAM tool are sent in the service even if it does not match the ethertype configured on the port on which the test SAP is configured. [195790-MI]

-
- 7210 SAS-K, CFM CCM packets not processed when the received CCM packet does not have End TLV. [195613-MI]
 - In case of 7210 SAS-D, recommended minimum value for EFM OAM timers is tx-interval of 500ms and a multiplier of 4. Even though CLI allows for lesser values, it is not supported - MI.
 - EFM-OAM sessions flap under the following conditions in case of 7210 SAS-D and 7210 SAS-E:
 - Using timers less than the default values,
 - STP packets that need to be forwarded in the slow-path (CPU-based forwarding) are incoming at a rate >64kbps, and
 - The CPU utilization is > 80% [76129-MA]
 - ICMP pings with higher packet sizes sent at higher rates will fail. [77611-MI]
 - The "vlan 0" for CLI "*configure>eth-cfm domain < > association < > bridge-identifier < >*" is not supported. [169536-MI]
 - Eth-CFM convergence fails; when llf is enabled on the LAG SAP and Lag shut and no-shut is done, or box is admin saved and rebooted. **Workaround:** Remove and add back MEPs. [116513-MA]
 - The ETH-CFM defect is not reported on booting the 7210 SAS with a configuration that has the ETH-CFM session over a LAG with no member ports. [100813-MI]
 - It is not recommended to use fault-propagation when the service entities (for example, SAP and SDP binding) are configured on a LAG. A LAG flap can result in CFM defect being raised which in turn results in a false fault propagation event. This issue is seen only when CCM timers of less than 10 seconds are in use. This configuration is not blocked in CLI but it is highly recommended that users do not use fault propagation with LAG when using CCM timers of less than 10 seconds. [150233-MA]
 - On 7210 SAS-D and 7210 SAS-T, on a given SAP in any service, in order to use Y.1731 2-DM, it is recommended to configure only one Y.1731 MEP, or configure MAC Address for each MEP created on the same SAP otherwise, 2-DM may not work properly. Except 2-DM, other functionalities are not affected even if there are multiple MEP's. [118015-MI]
 - For EFM OAM with timers less than the default values can result in EFM flaps due to events such as STP topology change, clearing FDB, or adding and removing ports to a LAG. It is recommended to use the default timers, tx-interval=10 and multiplier=5 or above. [81218, 82228-MI]
 - For 7210 SAS-E, under a heavy load condition, CFM/EFM/LACP/Eth-ring may flap due to any of the following triggers [74223, 106782]:
 - Executing a **clear fdb** or a **show fdb** command.
 - Transferring large files using FTP onto the flash.

- Mac moves occur at a rate of approximately 40 macs/sec.
- For 7210 SAS-E, CFM packets generated from the CPU are assigned to FC BE. Hence, there may be CFM flaps if the egress BE queue is congested.
- On PBB Epipe SAP, if the UP MEP is configured, it is recommended to enable CCM, with CCM disabled loopback, linktrace, DM, and SLM tests fails. [136751-MI]
- When FRR with facility backup kicks in and the merge point is the LSR with implicit null is 1 hop away, lsp-ping/trace does not work. [108025-MI]
- ETH-CFM Down MEP session on standby pseudowire does not come up in a VLL/VPLS pseudowire redundancy configuration. In case of standby becomes active ETH-CFM session comes up. [100594, 148081 -MI]
- LSP-trace fails over bypass tunnel when the LER is the PLR in FRR Facility backup configuration. [109165-MI]
- OAM "lsp-ping and trace" fails for ACH type "none" when configured on unnumbered-mpls-tp interfaces with "multicast/broadcast" static-arp. [180659-MI]
- For **mac-ping** and **mac-trace** when used with the **fc** and **return-control** options, do not use the egress queue as specified by the **fc** option. [82306-MI]
- In case of 7210 SAS-T, the Y.1564 testhead, when configured to use internal-loopback-ports, latency will not be computed when frame size is greater than 9000 bytes. Same test works fine with front panel port used as loopback ports. [177164-MI]
- For 7210 SAS-X, when port loopback with MAC swap is used, packets received on the test SAP (after going the test SAP loop) do not use SAP ingress queues. Only SAP ingress meters are used.
- Using **mac-populate** and **mac-purge** simultaneously on several VPLS services (for example, 64) could result in instability of the router. To avoid this anomaly, it is recommended to carry out this operation per VPLS service. [79690-MI]
- The timestamps are all 0s when **cpe-ping** is performed with SAA. [81726-MI]

11.11 QoS

- In case of 7210 SAS-K, in rare scenarios with non-default QoS policies, traffic issues may be seen for some Lag member links upon adding new link. Workaround is after adding new link, remove and re-attach Egress QoS policy, this workaround resolves the issue. [234245-MI]
- In case of 7210 SAS-R, 7210 SAS-T, 7210SAS-Mxp, and 7210 SAS-Sx, the Network Port Dot1p marking does not work for MPLS Packets. [195768-MI]

- For 7210 SAS-K MAC, IP and IPv6 based classification or ACL is configured, MAC/IP can occupy double wide slice. In scaled configuration attaching IPv6 classification/ACL may fail if there not enough entries in double wide slice. Number of available entries can be seen using "tools dump system-resources ingress-qos/ingress-acls". Workaround is to remove some of the MAC/IP policies, attach IPv6 policy. [199090-MI]
- 7210 SAS-K mac-criteria classification used on null or * SAP, null encap packets received are classified to either dot1p 0 or outer-tag 0 entry. If dot1p-classification policy is used then they are classified to default-fc. [216478]
- In 7210 SAS-R, RVPLS service the L2 unicast traffic is taking port based remarking, where as it is supposed to take sap-based remarking. [203154-MI]
- In 7210 SAS-R, if a LAG has more than 1 port from the same IMM then while attaching a queue-policy to the primary-port make sure the 'total CBS under that policy multiplied by the no. of ports in the lag' doesn't exceed max CBS value(140MB) available CBS value (140- MMU Configured CBS). [202780-MI]
- 7210 SAS-K, fragmented packets may be received upon changing the autoneg capability of remote device from 100Mbps to 1Gbps. [197188-MI]
- 7210 SAS-K, traffic having mixed frame sizes and also for jumbo frames deviation in shaper rates may be seen. [193169/195681-MI]
- 7210 SAS-K, classification based on inner Dot1p is not supported, though the CLI command match-inner-dot1p is available.
- 7210 SAS-K, a max of 200 queues for Ingress and 200 queues for Egress are supported. Allocations available in tools dump system-resources. The software does not enforce this check. It is not recommended to exceed this number at anytime.
- 7210 SAS-K, configuring Egress Rate less than 10Kbits/sec is not supported. [196558-MI]
- 7210 SAS-K, Max MBS for Queue supported is up to 11.5 MB even though configuration is available from CLI. [198927-MI]
- 7210 SAS-T network mode Dot1p remarking is not working for traffic ingressing I-SAP and egressing out of PBB B-SAP. [182552-MI]
- Ingress qos resources are not getting freed up when the ingress qos policy is removed from the SAP with tod-suite configured. [141896-MI]
- The maximum entry-id value for the IP/MAC-criteria in the SAP ingress policy context is restricted to 63 although the CLI allows up to 64. Entry-id 64 should not be specified. [76790/76964-MI]
- When remarking is enabled on access egress for a Dot1q port, the Dot1p bits in the customer tags get remarked, when the traffic is sent out of a Dot1q default SAP or a NULL SAP. [86818-MI]

- When mac-criteria dot1p-only is used as the classification criteria for a B-SAP and when resources are allocated from an IPv4 or IPv6 TCAM slice, the unlearned traffic from B-SAP to I-SAP hits the multicast meter even if an unknown meter is configured for that FC in the policy. This issue does not occur if the TCAM slice is of type MAC. [152477-MI]
- The slope policy attached to a hybrid port under the access-egress context will be ignored as it is unsupported. [161469-MI]
- L2PT tunneled STP packets do not have the appropriate MPLS EXP bits set. They are set to zero. [90580-MI]
- In case of 7210 SAS-R6, in rare occasions queue buffers not released during below operations while traffic flowing over SAP or network port, this can result in configuration failures and console messages.
 - attaching user defined queue policy
 - modification of queue scheduler mode, CBS, MBS values
 - port down
 - LAG port removal
 - SAP delete

It is recommended to carry out the above configuration changes after stopping traffic.

The configuration failures can be re-covered by rebooting line card.

In case of SAP delete, it is recommended to shutdown SAP before deletion. [168529-MI]

- In case of 7210 SAS-T, during heavy congestion in the system such that shared buffers are utilized completely and scheduler mode is strict, line rate traffic may not be achieved for all sizes of packet. [164543]
- In 7210 SAS-T access-uplink mode, inner tag dot1p not getting remarked when egress is qinq sap and ingress is dot1q sap. [163616]
- In 7210 SAS-X, it is not recommended to oversubscribe the 1.3GB buffer using ingress or egress queues. Oversubscription can result in traffic loss whenever new queues (ingress or egress) are created in HW due to events like enabling ingress queues on a SAP, adding a new link to a LAG SAP. [132888-MI]
- Ingress queue configurations is not supported through TOD. [149711, 149712-MI]
- In 7210 SAS-X, when scheduler-mode of the port is changed from sap-based to fc-based or vice versa, the traffic on the port may be disrupted for 2-3 seconds. [111643-MI]

- In 7210 SAS-X, when port scheduler mode is "fc-based" and multiple SAPs are configured on that port, and when one of the SAP consumes high bandwidth either due to queues to CIR levels, PIR weight, or CIR rate, remaining PIR traffic among rest of the SAPs having default SAP egress qos policy may not be fairly distributed. [107446-MI]
- In 7210 SAS-X, when egress rate command is used on network port, queues are not getting the expected rate for the following conditions:
 - When traffic flows from null sap to out of network port.
 - When queue rate is close to ERL rate. [109689-MI] [111255-MI]
- In 7210 SAS-X, during configuration oversubscription of queue CIR may result in undesirable behavior of shaper and scheduler. Oversubscription of CIR is not recommended. [101715-MI]
- The actual traffic rate may fluctuate from the configured CIR and PIR for the queue. The average rate is as per the operational rate but the rate may fluctuate around +/- 55-60kbps from the configured rate. [102303-MI]
- When the multiple queues have traffic flowing and the queues having same cir-level have very less pir bandwidth to share among them or when pir-weight ratio is very high for the queues under same cir-level, it is seen that traffic is not shared among same cir-level queues as per the configured pir-weight for the queues.
- Under a network port policy with Dot1p classification, traffic for classifiers with **out-profile** will be momentarily marked **in-profile** when some other classifier with profile **in** is changed to **out**. [85482-MI]
- In 7210 SAS-E, if a port on which egress-rate port limiting is enabled and is mirrored on the egress, then for packets less than 768 bytes, the rate of traffic seen on the mirror destination port is not the same as that on the mirrored source port. [83168]
- In 7210 SAS-E, when the queue on egress is within the CBS limit, the allocation of buffers is purely based on packet arrival time regardless of the packet's in-profile or out-profile state. This can sometimes cause unequal distribution of queue's CIR/PIR to the various ingressing traffic using the same queue. As a workaround, configure an aggressive slope policy for yellow packets. [74730]

11.12 Timing

- * Change in Slave MAC address is not updated under "show system ptp port detail", displays previous MAC address. [241905-MI]

- * In case 7210 SAS-R PTP G8275.1, when multiple links are present between master and slave, 7210 SAS-R may not select the correct master after a master loss event, It may choose a less preferable master. If there is single link between master and slave issue is not observed. [249786-MI]
- * In case 7210 SAS-R PTP G8275.1 1pps output error is not in-line with Time Error in the packets. [243483-MI]
- * PTP CLI's enabled for 7210 SAS-Sx Copper variant though PTP is not supported. [244321-MI]
- * parentIP field in "show>system>ptp" context will show the last master selected if the clock is in holdover. [252661-MI]
- * 7210 SAS-R, when the cpm switchover is performed, after the last active reference goes down, the new standby cpm should be in free-running state and after double switchover the dut should move to free-run state. In 9.0 release "show system sync-if-timing <standby>" shown as holdover state but clock is in free-run. [253442-MI]
- CLI "*configure system ptp profile "g8275dot1-2014"*" is not supported.
- NTP broadcast client server functionality not supported in RVPLS service. [240564-MI]
- 1 PPS output is not supported in case of 7210 SAS-R12 when PTP hybrid mode is configured. [221170-MI]
- In case of 7210 SAS-R, 7210 SAS-T, and 7210 SAS-Mxp, PTP hybrid mode with BC is deployed in a topology where a 2-way master/slave relationship exists, if master's time information changes drastically during reboot or initialization then sometimes few of them not to lock to the PTP master.
Note: This issue is not observed if all nodes in topology are 7210 nodes. Workaround for this issue is enable PTP only after its master is stabilized or reboot the nodes in a staged manner. [220749-MI]
- In case of 7210 SAS-T and 7210 SAS-Mxp, PTP hybrid mode is configured, ptp shut/no shut can result in 1-pps frequency not proper. Node reboot will resolve this issue. [220367-MA]
- 7210 SAS-K, PTP will only supported with singly tagged PTP frames with TPID=0x8100 and doubly tagged PTP frames with outer TPID=0x9100 and inner TPID=0x8100.[214723-MI]
- 7210 SAS-R, when a switch over from active to standby CPM is triggered, PTP on new active CPM takes about 40 minutes to go to a stable, locked state. The downstream slaves using this SAS-R observe a similar performance behavior. [210553, 201629-MA]
- "IEEE 1588/PTP Clock Recovery Event Statistics" of CLI "show system ptp statistics" gets reset after PTP shut/no shut. [150062-MI]

- With PTP configured, it is not recommended to remove system IP, this can result in PTP flaps. If system IP is shutdown when PTP uses system IP as source, the PTP sessions continue to use the system IP. [143255-MI]
- For 7210 SAS-X, when 1588 PTP slaves connected to upstream SASX as a 1588 BC, "Packet Loss" counter on slave increments every two minutes for CLI "show system ptp statistics". This does not affect frequency or time recovery. [147821-MI]
- For 7210 SAS-X, system deriving clock from 10 gig port with XFP (3HE00566CAAA01) does not move to holdover state when the master Dut is admin rebooted. Power cycle of the master Dut works fine. [105269-MA]
- For 7210 SAS-X, syncE not supported on the following SFP's.
 - Dual Rate SFP's: 3HE04116AA and 3HE04117AA
 - Fast Ethernet SFP's: 3HE00869AB, 3HE01454AAAA02 and 3HE00024AAAA02

11.13 Services

- * In case of 7210 SAS-K 2F4T6C, in some cases BGP 3107 session over RSVP with FRR configuration can result in node reboot, this configuration is not recommended with 7210 SAS 9.0R4 release. [252944-MA]
- * In case of 7210 SAS-K 2F4T6C, SAPs stats incrementing when it is in shutdown state. [252286-MI]
- * Dot1x is not supported on 7210 SAS-K 2F4T6C. [249623-MI]
- * Not all macs would get flushed when the service fdb size is changed on the fly on 7210 SAS-Sx. [251689-MI]
- In case of 7210 SAS-K 2F4T6C, if sdp-ping is done with a size equal to the service-mtu, the response fails, since 4 additional bytes are being sent. Workaround is to increase the service-mtu or use lesser value for "size" in the sdp-ping command.[MI-245001/246071].
- Status of control word in a service having MS PW is not propagated properly to the remote nodes. [241513-MI]
- MBB support for MVPN services is not available.
- In 7210 SAS-R, Eth-Ring switch over times are greater than 50ms with link failure which is farthest from the RPL node and more than 20 services protected by eth ring. [187004-MI]
- 7210 SAS-K, maximum 4 Access-Uplink ports are supported on the node, even though CLI allows configuration of 5 access-uplink ports.

- In 7210 SAS-R, STP state of saps on root bridge are incorrect after Clear Card and SWO in succession. [172199-MI]
- Dot1p priority is not preserved for RSTP packets forwarded by the node (for example doubly tagged STP packets received on a SAP). The dot1p is remarked to 7. Additionally, these packets are not matched against egress filter policies. [71855/75921-MI]
- When a time-range is expected to be active, a delay of up to 8 seconds can be expected for an associated filter policy on either ingress or egress. [161652-MI]
- Changing the **fdb-table-size** on multiple services simultaneously (for example, by a script) may affect MAC learning in some services. [82362-MI]
- There may be a small amount of traffic flow between services when LSPs are removed and added back within a very short duration, or when the LSPs are cleared under a given SDP. [84963-MI]
- On same Hybrid port if FRR for LSP and G8032 for SAP is configured, traffic switch over times observed for FRR and G8032 protected entities are high.[162771-MA]
- In case of PBB BCB (B-SAP to BSAP) packet, C-SA (customer mac) is learned in ivpls service when ivpls, isid matches transit packet's id.
Note: a small amount of C-SA (up to a maximum of few hundreds) are learnt and not all C-SA are learnt. [129525-MI]
- When SAS is configured as PE, in IGMP enabled VPLS service if spoke sdp is blocked due to reception of standby bit from connected MTU and then moves to forwarding state. There would be multicast data traffic loss till the next set of query/report is received. [121470-MI]
- PBB-packets received on B-SAP forwarded to epipe SAP if B-SA MAC matches "backbone-dest-mac" configured in pbb-epipe service. Packets with other B-SA MAC are dropped. [113910-MI]
- For SDP configured with vc-type vlan and "no vlan-vc-tag", tag value is carried as 0. In case vlan-vc-tag is configured, the proper tag is carried. [162708]
- In case of 7210 SAS-R6, service with L2PT/BPDU translation enabled, when CPM switchover happens, no xStp packets get translated in 4 seconds. Due to which xSTP flaps and a loop is created. [171060-MI]

11.14 Statistics/Accounting

- In case of SAS-R IMMv2, if terminated MPLS service packets has broadcast or multicast dest mac, then port ingress stats packet count does not increment. [201634-MI]

- In 7210 SAS-D and 7210 SAS-E, when * sap and other encap saps exist on same port, if egress statistics enabled on * sap and egress statistics disabled on the other saps, other saps egress packets are counted as part of * sap egress statistics. Similar behavior is observed with *.* or 0.* saps. [127452-MI]
- Attaching accounting-policy on SAS-D L2 Uplink SAP is not supported and should not be configured, however, CLI commands exists. [111897-MI]

11.15 STP

- If a large number of MAC addresses exists in the VPLS FDB and the entire FDB is flushed and relearned, there may be a period of when RSTP BPDUs are not sent. A partial workaround is to configure fdb-table-size limits. [40532-MI]
- For 7210 SAS-E, if an ingress CPU forwarded RSTP/L2PT/PVST packet rate exceeds 64kbps, RSTP flaps. On SAS-M, if an ingress CPU forwarded RSTP/L2PT/PVST packet rate exceeds 128kbps on access, or 200kbps on network port, RSTP may flap. [73189]
- If 7210 SAS-Ms are connected in a ring topology with LAG configuration and FRR, when ring is broken, STP flaps for some service is observed. [95969-MI]
- RSTP convergence fails if **force-vlan-vc** is enabled on a mesh SDP. [94928-MA]

11.16 System

- In case of 7210 SAS-K combo port does not detect dualRate sfp (3HE05164AA/3HE09500AA) in "Connection-type Auto".[240946-MI]
- 7210 SAS-Mxp,7210 SAS-T, 7210 SAS-D, DDM read failure messages seen with 1 gig (GIGE-SX) SFP vendor type OPC 3HE00027AA.[216028-MI]
- 7210 SAS-Mxp, Copper SFP is not supported on port 23 and 24.
- 7210 SAS-K, 100FX SFPs are not supported.
- 7210 SAS-K, Copper SFP is supported only in 1Gig mode in port 1 and 2.
- 7210 SAS-K, Copper SFP is not supported on port 3.
- In 7210 SAS-D, if port is not shut before loopback is configured, speed is always set to auto negotiated speed. If port is shut before loopback then speed set to 1Gbps. [114904-MI]
- 7210 SAS-D, sometimes "[iomMsg-1]soc_phyctrl_loopback_set: u=0 p=3 TIMEOUT" console message appears when port internal loopback is enabled/disabled on Copper SFP port. [115784-MI]

- 7210 SAS-D, snmpwalk reports redundant power-supply status. [113827-MI]
 - In certain scenarios, traffic drops are seen during source learning when copper ports are operating in half duplex modes. Once learned, there are no other observed drops. [75875-MI]
 - The system time MIB object stiDateAndTime is the UTC time and should not include the time zone offset in SNMP **get** and **set** requests. [66553-MI]
 - During auto-config, if a DHCP relay packet for an unrelated DHCP session is received by the system, the system may be non-responsive. [76811-MI].
 - When using 100FX SFPs, autonegotiation should be disabled. [99132-MI]
 - The system saves core dumps at the URL specified in the BOF, while coming up. Specifying a remote location as a core file destination has two issues.
 - If a remote URL is specified, and if the uplinkA cannot be brought up, saving the core dump may not succeed, even if uplinkB is up.
 - If uplinkA's BOF is configured to use DHCP, the IP address acquired by DHCP is not released by the system after saving the core file.
- Thus, it is recommended to configure the core file destination to be the local flash (cf1). [76764, 76736-MI].
- Internal loopback of port does not work when loopbacked port speed is set to 100Mb and autoneg is turned off. [113995-MI]
 - In certain scenarios (such as Cu SFP connected to Fixed Cu), the show port command does not display the MDI/MDX value as expected. Functionality is not affected. [76765-MI]
 - During auto-config, if a DHCP relay packet for an unrelated DHCP session is received by the system, the system may be non-responsive. [76811].
 - In case 7210 SAS-E, M, D, applying large router policy configurations (2000 to 3000+ lines) at one time, can cause protocols such as LACP, EFM, to flap. [83787]
 - The statistics and the utilization rate displayed by the **monitor port rate** CLI command for a given time interval does not match the actual count and rate received by the system in that time interval. [83757]
 - During Autolnit using DHCP, it is recommended to store proper bof, config files in defined location. In lab scenario, it is observed that, through DHCP if 7210 SAS fetches the ip address, but bof or config file is not stored in defined location, after several DHCP request or replies, DHCP starts failing. A reboot of 7210 SAS is required to recover from DHCP failures. [131915]
 - Control protocols may flap when the **file copy** command is initiated, using FTP, from the system if the specified FTP server is not reachable. This happens under loaded CPU conditions. Before initiating the **file copy** command using FTP, verify that the FTP server is reachable. [76566].

-
- In a scaled configuration with STP enabled, when on an event (such as a network LAG flap) causes an enormous amount of traps sent out of the 7210, STP may flap. [95969-MI]
 - 10Gig alarms, “no-frame-lock” and “high-ber”, are not supported on the 7210 although they are configurable.
 - The **configure port *port-id* ethernet report-alarm** command option **no-frame-lock** or **high-ber** are allowed to be configured but they are not supported. This command is applicable only for 10G ports.
 - In loaded CPU conditions and scaled configurations, CPU spikes may cause STP to flap. This was seen in one condition when CPU spiked due to link down event and was processing 16K OSPF routes. [84889-MA]
 - In loaded CPU conditions and scaled configurations, CPU spikes may cause LAGs to flap. This was seen under the following events:
 - Addition and deletion of an OSPF interface in non-backbone area. [84537-MI]
 - Addition and deletion of port (or LAG port) to ip-interface. [84165/86007-MI]
 - Enabling an SSH server. [84166-MI]The workaround or to lessen the chances of flaps occurring in this case, it is suggested to enable the **preserve-key** option before enabling SSH.

11.17 Hardware

- Only one port LEDs is used to indicate port status and link activity. If the LED is lit steady, it indicates link up. If the LED is blinking, it indicates link activity.

Customer Document and Product Support



Customer documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation feedback

[Customer Documentation Feedback](#)

