Alcatel-Lucent

Service Access Switch| Release 8.0 Rev.02

7210 SAS OS Software Release Notes

3HE10395AAABTQZZA V8.0.R2 Issue 1

Alcatel·Lucent

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

Table of Contents

RELEASE NOTE ORGANIZATION FOR 7210 SAS-D, E, K, M, X, T, MXP, AND R6

The following are the major topics covered in these Release Notes:

RELEASE 7210 SAS SUPPORTED HARDWARE

- 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

SOFTWARE UPGRADE NOTES

UPGRADE TO 8.0R1 OR LATER RELEASES

- In 8.0R1 release following accounting numbers are modified, new numbers are shown in( ) netInfIngressOct(101), netInfIngressPkt(102), combinedNetInfIngress(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.

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 "c*onfig>system>res-profile>router*".

UPGRADE TO 7.0R1OR LATER RELEASES

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 7.0R1or 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.

UPGRADE TO 6.0R6 OR LATER RELEASES

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.

UPGRADE FROM 6.0R4 OR PRIOR RELEASES

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.

UPGRADE FROM 3.0 RELEASES

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-expprofile- 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.

UPGRADE FROM 4.0 OR 5.0 RELEASES

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.

UPGRADE FROM 5.0R1 OR 5.0R2 RELEASES

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.

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
```

UPGRADE TO 5.0R1 OR LATER RELEASES FROM PRIOR RELEASES

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 ( ) netInfIngressOct(52), netInfIngressPkt(53), combinedNetInfIngress(54), accessEgressPkt(55), accessEgressOct(56), combinedAccessEgress(57), combinedNetEgress(58), combinedSvcEgress(59), combinedSvcInEgPkt(60), combinedNetInEgPkt(61)

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.

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

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

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.

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.

UPGRADE TO 4.0R4 OR LATER RELEASE FROM PRIOR RELEASES

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.

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.

UPGRADE TO 1.1R7 OR LATER RELEASES FROM PRIOR RELEASES

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>boot-bad-exec “cf1:/<config-file>”

Step 3. Follow the Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, and Mxp devices on page 10.

UPGRADE TO 1.1R6 OR LATER RELEASES FROM PRIOR RELEASES

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>boot-bad-exec “cf1:/<config-file>”

Step 3. Follow the Software Upgrade Procedures for 7210 SAS-D, E, K, M, X,T, and Mxp devices on page 10.

SOFTWARE UPGRADE PROCEDURES FOR 7210 SAS-D, E, K, M, X,T, AND MXP DEVICES

The following sections contain information for upgrading to the 8.0R1 software version. In particular, there are sections that describe the following:

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

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:\8.0R2 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.

Note: Applicable only to 7210 SAS-T, 7210 SAS-K and 7210 SAS-Mxp devices

- The 7210 SAS-T, 7210 SAS-K, 7210 SAS-Mxp 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 6.0R6, 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.

SOFTWARE UPGRADE PROCEDURES FOR 7210 SAS-R6

The following sections contain information for upgrading to the 8.0R1 software version. In particular, there are sections that describe the following:

- Standard Software Upgrade Procedure

  Procedure for performing a standard, service-affecting upgrade.

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-R6 OS Images to cf2:

The 7210 SAS-R6 image files must be copied to the cf2: device on the 7210 SAS-R6 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:\8.0R2.

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-R6 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 7.0R4, 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 Alcatel-Lucent 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.

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

RESOLVED IN R8.0R2

The following are specific technical issues that have been resolved in Release 8.0R2 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp and 7210 SAS-R6 OS.

- 7210 SAS-Mxp non-ETR platform is supported from 7210 SAS 8.0R2 release. [218331-MA]
- In previous releases, in the case of 7210 SAS-M access-uplink mode, 7210 SAS-T access-uplink mode and 7210 SAS-E, filter attached to LAG SAP was not effective for traffic entering member port of lag, where member port is from 1/1/12 onwards. Workaround was to remove and add back filter policy to LAG SAP. This issue is resolved in 7210 SAS 8.0R2. [215495-MI]
- In previous release, IPv6 filters attached to VPRN SAP (dot1q encap and port mode hybrid) at ingress stopped working after reboot with saved configuration. Workaround was to remove and add back filter policy. This issue is resolved in 7210 SAS 8.0R2. [214924-MI]
- In previous release, in 7210 SAS-K, mac-filter, ip-filter, ipv6-filter, and sap-ingress policy id configured should be between 1 to 1000, this issue is resolved in 7210 SAS 8.0R2. [215580]
- In previous release, with IGMP snooping enabled, multicast traffic getting flooded to all ports in the network lag over which the sdp is formed in case of 7210 SAS-Mxp. The same issue was observed on 7210 SAS-R6, when LAG ports are on same IMM. This issue is resolved in 7210 SAS 8.0R2. [MI – 213380].
- In 8.0R1 release of 7210 SAS-K, with default static-route configured (0.0.0.0/0), IP packet with destination IP as the system IP or loopback address were not processed by system, instead they were forwarded using default route. The workaround was not to configure default static-route when system interface / loopback interface configured. This issue is resolved in 7210 SAS 8.0R2. [216401-MA]
- In previous release, Y.1731 version-1 DM request packets received on UP MEPs were not processed on 7210 SAS-R6. This issue is resolved in 7210 SAS 8.0R2. [216007-MI]
- In previous releases of 7210 SAS-K, in sap-ingress qos policy, 'mac-criteria dot1p-only', 'ip-criteria dscp-only' and 'ipv6-criteria dscp-only' configuration though CLI was allowed, from 8.0R2 these are removed from the CLI. Always use the default type. [216112-MI]

- In previous release, 7210 SAS-K un-tagged dot1x packets were not being tunneled with or without enabling tunneling under the port in both VPLS and Epipe services, this issue is resolved in 7210 SAS 8.0R2. [216225-MI]
- In previous releases, SAM sync performed on 7210 was failing if there was MTU mismatch in the SAS to SAM path and node is managed over IPv6 in-band. This issue is resolved in 7210 SAS 8.0R2. [214627-MA]
- With 7.0 release when 7210 SAS-T devices, deployed in a ring topology and PTP is configured in hybrid mode with BC, parallel reboot of all the 7210 SAS-T ring nodes sometimes used to cause few of them not to lock to the PTP master. Workaround for this issue was to reboot the 7210 SAS-T ring nodes in a staged manner. This issue is resolved in 7210 SAS 8.0R1. [209601-MA]
- Log messages relating to G.8032 Ethernet Ring Protection Switching have been modified to reference paths a/b instead of paths 0/1. [215300-MI]
- In previous release of 7210 SAS-Mxp, due to KeepAlive failure some of the SDP's are flapping if SDP's are scaled number. This issue is resolved in 7210 SAS 8.0R2. [215582-MI]
- In previous release of 7210 SAS-K, if dot1x enabled on port before adding port to LAG, then dot1x packets were not processed by 7210 SAS-K. Same issue happens when port removed from LAG. This issue is resolved in 7210 SAS 8.0R2. [216125-MI]

RESOLVED IN R8.0R1

The following are specific technical issues that have been resolved in Release 8.0R1 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp and 7210 SAS-R6 OS.

- In 8.0R1 Eth-ring control SAP configuration in RVPLS service is not allowed, In previous release such configuration on 7210 SAS-M network mode, 7210 SAS-X, and 7210 SAS-R6 reboots the node. [210392-MA]
- In previous release, "configure mirror no mirror-dest <>" results in following messages on console for SAS-R6 IMMv1 " 1:iomMsg-1:IOM:is_group_empty Getting Group Info Failed : Entry not found". This is resolved. [199727-MI]
- On SAS-R6, Epipe MIPS require resource to be configured for "down-mep" and "up-mep" under "configure system resource-profile", from 8.0R1 MIP creation not allowed if resource not allocated. [202895-MI]
- In previous release, In case of UP MEP, until peer MAC is learnt in VPLS service all CFM packets were sent with dot1p value "0", this is corrected now.[201342-MI]
- In previous release, 7210 SAS-R6 UP MEP CCM messages were sent with dot1p 0 in VPLS service when UP MEPs are configured on SAPs and egress is SAP, this is corrected now. [201177-MI]
- In previous release, on 7210 SAS-K, CFM 2-DM tests did not work when disable-learning enabled in a VPLS service, this issue is resolved. [198054-MI]
- Under "config>system>thresholds#" CLI context "kb-memory-use-alarm" and "kb-memory-use-warn" commands were missing in earlier releases. These commands are available now. [210697-MI]

- In previous releases, packet rate and utilization shown in “monitor port x/y/z rate” command are not accurate if the polling interval is non-multiple of 10sec. In this release, accuracy of the rates is improved. [211466-MI]
- On SAS-X, if the number of G.8032 rings configured are greater than 22 and the system is upgraded to release 6.0/7.0, system crashes. This issue is resolved in this release. [212517-MA]
- On SAS-R6, “show pool” output was showing incorrect CBS and MBS values if the system is brought up in port-scheduler mode. This is resolved in this release. [213033-MI]
- In previous releases When down MIPs are present in both ingress end and egress end in VPLS service, the MIP on egress end also used to respond to LT requests though it is only down MIP. This error has been rectified, from 8.0R1 down MIPs respond only in ingress direction irrespective of presence of MIP at egress end in VPLS service. [190115-MI]
- In earlier release of 7210 SAS-K booting through port 3 is supported in fiber mode (SFP) only, copper port can also be used from 8.0R1. [202510-MI]

NEW FEATURES

RELEASE 8.0R2

The following items describe features added to 8.0R2 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp and 7210 SAS-R6 OS.

PER SAP EGRESS RATE-LIMIT OPTION USING SAP EGRESS METERS

Platforms Supported: 7210 SAS-D, 7210 SAS-M and 7210 SAS-T (in both access-uplink and network mode)

In prior releases, the platforms listed above (that is, 7210 SAS-D, 7210 SAS-M, and 7210 SAS-T) supported port-based egress queues with an option to limit of traffic per FC or per port, but with no option to limit the amount of traffic sent out a SAP. With this release, a meter can be associated with SAP egress to limit the aggregate amount of traffic (across all FCs) sent out of the SAP. This feature is useful to limit the amount of traffic sent out of a SAP in a P2MP service (For example: VPLS Service). This feature is supported for SAPs in Epipe, VPLS, RVPLS, IES and VPRN services.

NOTES:

- Before configuring egress aggregate meter on SAP, user must reallocate the resources from the egress-internal-tcam pool for use with this feature using the CLI command "*configure> system> resource-profile> egress-internal-tcam> qos-sap-egress-resource*". The egress internal- tcam pool resources are shared with other features such as egress ACLs. User needs to take away resources from those features and reallocate them towards SAP egress meters.
- The per SAP egress aggregate policer is not service FC and profile aware. It looks at the packets scheduled out of the per port queues on a first-come-first-basis and limits the amount of traffic across all the FCs.

For more information, see 7210 SAS-D, E, K Basic System User Guide, 7210 SAS-M, T, X, R6, Mxp Basic System User Guide, 7210 SAS-D, E, K Services User Guide and 7210 SAS-M, T, Mxp Services User Guide.

RELEASE 8.0R1

The following items describe features added to 8.0R1 of the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-X, 7210 SAS-T, 7210 SAS-Mxp and 7210 SAS-R6 OS.

7210 SAS-MXP PLATFORM

Release 8.0R1 introduces support for the new platform - 7210 SAS-Mxp 22F 2C 4SFP+. 7210 SAS-Mxp can be used as 1G/10G access aggregation or 10G Demarcation device extending IP/MPLS to the edge of the access network. 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 SAP hierarchical egress shaping/scheduling, along with extensive Ethernet and MPLS OAM support. It provides support for network synchronization capabilities such as syncE and 1588v2/PTP.

7210 SAS-Mxp is a 1.5RU, 19-inch rack mountable, NEBS and ETSI compliant unit, available in 2 variants:

- 7210 SAS-Mxp - Supports standard Temperature range of 0oC - 40oC
- 7210 SAS-Mxp ETR – Supports extended temperature range of -40oC to +65.
- Provides rack and wall mount options.

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

- 22 x 100/1000 1G SFP ports, 2 Combo ports (SFP and RJ.5 Copper) and 4 x 10G SFP+ ports
- 64Gbps (Full-duplex) of switching capacity with IMIX traffic
- Redundant hot-swappable power supply (AC, DC -48V, DC +24V), with DC power source failure detection
- Hot-swappable fan tray with 3 fans and fan filter
- 2 alarm-output pin and 4 alarm-input pins, with an option to supply power (+24V) to the alarm-input.
- 3 storage locations – Internal non-replaceable flash (cf1:\) of size 2GB, external field replaceable flash (cf2:\), and USB (uf1:\) storage device.
- 7210 SAS-Mxp ETR units support Power over Ethernet (PoE) (802.3af) and PoE+ (802.3at) with the combo ports when it is configured for use as copper port. A maximum of up to 60W of power is available for use by connected PoE devices.

NOTES

- 7210 SAS-Mxp ETR requires use of 200W power supplies.
- 7210 SAS-Mxp can operate only in MPLS mode (also known as, network mode). By default, the box boots up in network mode and no user configuration is needed to use MPLS functionality.

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

- Support Access, Network and Hybrid port mode
- Support the following SAP encapsulations – NULL, Dot1q, Dot1q Explicit NULL, Dot1q Default, Dot1q range SAP, QinQ (including Q1.*, Q1.Q2 and 0.* SAP)

- Service Support.
  - Epipe, VPLS, IES and VPRN services
  - Epipe services with support for both T-PE and S-PE along with Multi-segment PW and BGP PW routing
  - VPLS service, with support for BGP-AD (auto-discovery)
  - L2PT and BPDU tunneling in VPLS service
  - IGMP (v1 and v2 only) snooping (Layer-2 multicast) and MVR
  - IGMPv3 snooping is not supported
  - DHCP snooping over only SAPs
  - Layer-2 control protocol tunneling support for EFM, LLDP, 802.1x, and LACP in both Epipe and VPLS services
  - IES IPv4 services with support for OSPFv2, IS-IS, static routing and VRRP
  - VPRN IPv4 services with support for eBGP, OSPFv2, and static as PE-CE routing protocols
  - RVPLS service is not supported in this release
- MPLS Support
  - MPLS support for Epipe, VPLS and VPRN services
  - RSVP-TE with FRR (one-to-one and facility with PHP), primary & secondary LSPs with hot standby, SRLG, admin-groups, and others.
  - LDP with ECMP support for LSR LSPs
  - LDP DoD (server only) support
  - LDP over RSVP
  - BGP 3107 labeled routes for only L2 VPN and L3 VPN services with an option to install only those routes which are required for services configured on the node (a.k.a. BGP 3107 optimization).

    NOTE: For BGP 3107 and LDP-over-RSVP only FRR one-to-one is supported.
- IPv4 Routing Support
  - IPv4 forwarding support with static routing
  - OSPFv2 (single instance) and IS-IS (multi-instance)
  - VRRP support for IPv4 interfaces in IES IPv4 and VPRNv4 services
  - In BGP address family, only vpn-ipv4, ipv4-labeled routes (BGP RFC3107), IPv4 (only for PE-CE routing in VPRNv4 services) and l2-vpn families are supported. BGP IPv4 family is not supported in the base routing instance.
  - Supports route policies for management and control of distribution of routing information
  - DHCP (IPv4) relay support for IES and VPRN services
  - IPv4 multicast with support for PIM-SM, PIM-SSM and IGMPv3. Support is compatible with support on 7210 SAS-M/X.
  - BFD - This release supports BFD (both single hop and multi-hop) on IPv4 IP interfaces. BFD is supported in hardware and supports 10ms minimum timers for BFD sessions for IPv4 IP interfaces configured on a port. For BFD sessions for IPv4

IP interfaces configured on a LAG or using the system IP interface/loopback addresses, CPM-based sessions with a minimum timer of 100ms. BFD support is available for use with following:

- Static routes
- OSPFv2
- IS-IS for IPv4
- VRRP for IPv4 on IES, VPRN and network ports
- RSVP-T
- TLDP
- Multi-hop BFD for MP-iBGP sessions

- QOS and ACL SupportE
  - 8 Forwarding classes
  - SAP ingress QoS with hierarchical policing (2 levels - per FC, per SAP)
  - SAP egress QoS with 8 queues per SAP, with hierarchical shaping (3 levels - per FC, per SAP egress shaper and per port egress shaper). SAP based scheduling for SAPs, with support for Strict-priority scheduling (SP) and WDRR.
  - Network IP interface and network port ingress policing (only per FC)
  - Network port and Hybrid port egress queues with 8 queues per port, per port shaping and scheduling. Port based scheduling for hybrid ports and network ports, with support for Strict-priority scheduling (SP) and WDRR.
  - Around 145MB of packet buffers for burst absorption. The buffers are shared among SAP egress queues and network port egress queues. CBS and MBS parameters are configurable per queue.
  - WRED mechanisms with 2 slopes (high-priority and low-priority) per queue for congestion management.
  - SAP ingress classification supports MAC criteria, IPv4 criteria, and IPv6 criteria (as applicable)
  - Network ingress classification support MPLS EXP or IP DSCP and Dot1p priority bits.
  - Egress marking using Dot1p, IPv4 DSCP and MPLS EXP (as applicable) [NOTE: unlike 7210 SAS-M, 7210 SAS-Mxp, requires use of remarking policies for configuring egress FC to priority bit marking).
  - Supports DEI classification and marking
  - 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
  - Supports Synchronous Ethernet on SFP ports, SFP+ ports and combo ports when the connection-type is copper or SFP (only Fiber SFP).
  - User must use SFPs and SFP+ that support syncE to make use of the syncE support
  - Combo ports configured in copper mode, can be used to distribute frequency (master) or recover frequency (slave)
  - BITS, interface is supported
  - 1pps, ToD and 10MHz interfaces are not supported in this release
- High Availability and Reliability Support
  - Hot-swappable Redundant Power supplies
  - Hot-swappable Fan tray with 3 fans, with notification for a single fan failure
  - LAG with active/active and active/standby support
  - MC-LAG support (server support) with and without LACP
  - G8032 with the capability to use 7210 SAS-Mxp as interconnection nodes in a major-ring/sub-ring topology
  - STP, RSTP, MSTP with mVPLS/xSTP support
  - MPLS FRR - facility with 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
  - VRRP (IPv4) support in IES and VPRN services
  - Fault propagation support in Epipe service (For example: LLF, and others.)
  - BFD support with 10ms timers for faster failure detection (as a BETA only in 8.0R1)
- OAM support
  - Supports EFM OAM with support for EFM OAM dying gasp message or a SNMP dying gasp message on loss of power.
  - Supports LLDP
  - CFM/Y.1731 Down MEP, UP MEP, Ingress MIP only for VPLS, Ingress and Egress MIP for Epipe (see user guide for MEP support per service and different service objects).
  - CFM and Y.1731 based OAM tools – Supports CCM, Linktrace, Loopback, 2-DM, 1-DM, 2-SLM, AIS, and RDI.
  - MPLS OAM tools for Epipe, VPLS and VPRN services (example – lsp-ping, vccvping, vprn-ping, mac-ping, and others.).
  - Mirroring support – Only port ingress and egress mirroring with the capability to use NULL as mirror destination is supported in this release.
  - TWAMP
  - Ethernet CRC error monitoring
  - 1830 VWM device management for CWDM devices is supported

- Accounting, Security and Management support
  - Per SAP ingress and egress accounting records
  - Per network IP interface and network port accounting records
  - Out-of-band Ethernet management port is available with IPv4 support
  - Support for Dot1x is available
  - SNMP (including v3 support), SSH, Telnet, NTP, and others are supported
  - RADIUS & 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. 7210 SAS-Mxp provides the following boot options:
    - Using Autoinit
    - Using internal flash (cf1:\)
    - Using external flash (cf2:\)
    - Using USB port (uf1:\)

When shipped from factory the device is configured to use autoinit by default. For more information reference on how to use the various options to boot the 7210 SAS-T device to the 7210 SAS-Mxp Installation Guide and 7210 SAS-Mxp Basics System Configuration Guide.

Some of the features that are not supported in this release are:

- PBB is not supported;
- MPLS-TP is not supported;
- IPv6 for services and management is not supported;
- RVPLS Services is not supported;
- Remote Mirroring using Dot1q SAP and or SDP is not supported;
- Port loopback with MAC swap
- Y.1564 testhead OAM tool for service performance measurement before service turn-up
- PTP is not supported

For more information about the features and functionality refer to the 7210 SAS-M, T, Mxp User Guide set.

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

Platforms Supported: 7210 SAS-Mxp and 7210 SAS-T (network 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 discovey. It is not recommended to configure 7210 SAS device 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 l3-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, R6, Mxp Routing Protocols User guide, 7210 SAS-M, T, Mxp Services guide, and 7210 SAS-M, T, Mxp QoS guide for more information.

BGP ADD-PATH AND BGP PIC (PREFIX INDEPENDENT CONVERGENCE) FOR MPLS-BGP IPV4 AND IPV6 L3VPN ROUTES

Platforms Supported: 7210 SAS-R6, 7210 SAS-M (network mode), 7210 SAS-X

BGP Add-Path capability allows a BGP router to send to its BGP peers more than one path for the same prefix/NLRI. The base BGP standard does not provide for such a capability; if a BGP router learns multiple paths for the same NLRI, it selects only one as its best path, and this is the only path that is allowed to be advertised to BGP peers. Advertising multiple paths has several benefits including reduced routing churn and faster convergence (7210 SAS does not support BGP multipath for load sharing of traffic over multiple BGP path). The ability to send and/or receive multiple paths per prefix/NLRI must be negotiated with a BGP peer during capability negotiation. This release supports Add-Path for MP-iBGP sessions established by the VPN-IPV4 and VPN-IPV6 routes. It is not supported for eBGP sessions (for both IPV4 and IPv6) and for VPRN PE-CE BGP sessions, LBL-IPv4, IPv6, VPN-IPv6 and 6PE routes. The maximum number of paths to send is configurable per peer and per address family. The actual set of advertised paths depends on the number available in the RIB-IN, the next-hop diversity of the paths, BGP route advertisement rules, export policies and other BGP configuration options.

BGP fast reroute introduces the concept of a backup path for BGP routes to enable faster convergence in case of network failure. When multiple paths are received for an IP prefix/NLRI, usually only the best path(s) are downloaded to the data-plane for installation in

the IP FIB. When failures occur and all of the best paths have become invalid, failover to the next-best path takes time because the BGP decision process and FIB update must be executed for every affected prefix. When BGP fast reroute is enabled, a diverse next-hop backup path is provided along with all of the best/primary paths when the CPM downloads BGP routes to the data-plane. This extra information allows the data-plane software to group routes having the same order of preference of primary next-hops and backup next-hop so that they can reference a common route destination in the IP FIB. This grouping allows for a failover time that is prefix-independent (not dependent on the number of grouped prefixes). When all of the primary next-hops in a route destination have failed, the data-plane software re-programs the route destination to point to the backup next-hop. Since this is done without involvement of the BGP control plane, the failover is faster once the failure has been detected and signaled to the data-plane. This feature is only supported for MP-BGP IPv4 and IPv6 L3VPN routes. It is not supported for BGP 3107 labeled routes or any other routes advertised using BGP.

For more information, please see the 7210 SAS-M,X,T,R6,Mxp Routing Protocols User Guide.

OSPF Multi-Instance

Platforms Supported: 7210 SAS-R6.

On 7210 SAS-R6, support to instantiate more than one instance of OSPFv2 (IPv4 only) has been added. It allows for configuration of smaller network routing domains such that routing information can be separated in a large network with large number of nodes. Smaller routing domains allows the use of network devices with smaller IPv4 FIB capacity to be used in network deployments.

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

LDP FRR and IP LFA (Loop Free Alternate)

Platforms Supported: 7210 SAS-R6, 7210 SAS-X, 7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-M (network mode)

LDP Fast Re-Route (FRR) allows the user to provide local protection for an LDP FEC by precomputing and downloading both a primary and a backup Next-Hop Label Forwarding Entry (NHLFE) for this FEC to the data-plane.

The primary NHLFE corresponds to the label of the FEC received from the primary next-hop as per standard LDP resolution of the FEC prefix in RTM. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop. The LFA next-hop pre-computation by IGP is described in RFC 5286 – “Basic Specification for IP Fast Reroute: Loop-Free Alternates”. LDP-FRR relies on using the label-FEC binding received from the LFA next-hop to forward traffic for a given prefix as soon as the primary nexthop is not available. This means that a node resumes forwarding LDP packets to a destination prefix without waiting for the route convergence. The label-FEC binding is received from the loop-free alternate next-hop ahead of time and is stored in the Label Information Base since LDP on the SR OS operates in the liberal retention mode.

This feature requires IGP to perform the Shortest Path First (SPF) computation of an LFA nexthop, in addition to the primary next-hop, for all prefixes used by LDP to resolve FECs. The IGP also populates both routes in the Routing Table Manager (RTM). The IGP supported for LDP-FRR in this release is OSPFv2 and IS-IS.

NOTE: IP FRR is not supported for both IPv4 and IPv6 routes. Only IP Loop Free Alternate (LFA) is supported for IPv4 routes for use with LDP FRR.

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

OPTION TO REALLOCATE NETWORK QOS CLASSIFICATION RESOURCES

Platforms Supported: 7210 SAS-R6.

With this release, option is provided to reallocate resources allocated by default from the ingress-internal-tcam resource pool towards network port ingress and network IP interface ingress QoS classification. It provides flexibility to use the classification resources towards other features, such as SAP ingress classification, Ingress ACLs, etc. which share the resources available in the ingress-internal-tcam resource pool.

If resources allocated towards network classification are reallocated then software will error out an attempt to change the mode of the port to either network or hybrid mode. In other words, with no resources allocated for network classification, only access ports with access SAPs can be configured on the node. User has an option to reallocate resources allocated towards network classification only on some IMMs and retain the default allocation on some other IMMs.

NOTE:

1. Option is available to reallocate all the 2 chunks of resources allocated by default towards network classification. In other words, partial reallocation is not supported.

2. By default, when the IMM comes up, all the ports will be in network mode. If resources allocated towards network classification are to be reallocated then the mode of all the ports on that IMM should be changed to access mode else software will error out.

For more information, see 7210 SAS-M,T,X,R6,Mxp Basic System User guide and 7210 SAS-X,R6 QoS user Guide.

PER SAP EGRESS RATE-LIMIT OPTION USING SAP EGRESS METERS (BETA ONLY)

Platforms Supported: 7210 SAS-D, 7210 SAS-M and 7210 SAS-T (in both access-uplink and network mode)

In prior releases, the platforms listed above (i.e. 7210 SAS-D, 7210 SAS-M, and 7210 SAS-T) supported port-based egress queues with an option to limit of traffic per FC or per port, but with no option to limit the amount of traffic sent out a SAP. With this release, a meter can be associated with SAP egress to limit the aggregate amount of traffic (across all FCs) sent out of the SAP. This feature is useful to limit the amount of traffic sent out of a SAP in a P2MP service (e.g. VPLS Service). This feature is supported for SAPs in Epipe, VPLS, RVPLS, IES and VPRN services.

NOTE:

1. Before configuring egress aggregate meter on SAP, user must reallocate the resources from the egress-internal-tcam pool for use with this feature using the CLI command "configure> system> resource-profile> egress-internal-tcam> qos-sap-egress-resource". The egress-internal-tcam pool resources are shared with other features such as egress ACLs. User needs to take away resources from those features and reallocate them towards SAP egress meters.

2. The per SAP egress aggregate policer is not service FC and profile aware. It looks at the packets scheduled out of the per port queues on a first-come-first-basis and limits the amount of traffic across all the FCs.

For more information, see 7210 SAS-D,E,K Basic System User guide, 7210 SAS-M,T,X,R6,Mxp Basic System User guide, 7210 SAS-D,E,K Services user guide and 7210 SAS-M,T,Mxp Services user Guide.

MIPs (both Ingress and Egress) with primary VLAN support for VPLS services

Platforms Supported: 7210 SAS-R6 IMMv2 only

In prior releases, only Ingress MIPs (Maintenance Intermediate Point) were supported in VPLS service for both SAPs and SDPs. Ingress MIP responds only to messages that ingress the node. With release, MIP support is enhanced to support MIP both ingress and egress direction is available for use with SAPs in a VPLS service. MIPs are useful to operators to diagnose and localize faults by using CFM linktrace and loopback messages.

In addition, support has been added for use of primary VLAN with MIPs. This allows operator to use ETH-CFM tools for fault diagnosis and troubleshooting services when a SAP aggregates a set of VLANs (e.g. Dot1q Default SAP or Q1.* SAP). It allows the user to specify the VLAN which should be used to extract and process CFM messages in the context of the MIP (rather than using the default value – which is CFM messages received with no tags).

NOTE: Before using this feature user must reallocate the resources from the egress-internal-tcam pool for use with this feature using the CLI command configure> system> resource-profile> egress-internal-tcam> eth-cfm> bidir-mip-egress. The egress-internal-tcam pool resources are shared with other features such as egress ACLs. User needs to take away resources from those features and reallocate them towards this feature. In addition, the CLI command config>service>vpls>eth-cfm> vpls-sap-bidir must be enabled to use this feature.

For more information, see the 7210 SAS-X,R6 Services User Guide and 7210 SAS-M,T,X,R6,Mxp OAM and Diagnostics User Guide.

Support for xSTP and mVPLS/xSTP

Platforms Supported: 7210 SAS-K

This release adds support for STP, RSTP and MSTP on 7210 SAS-K providing another option for L2 resiliency and loop prevention mechanism. It also includes support for mVPLS (Management VPLS) which provides an option to use a single control instance of the xSTP protocol and use the state of the control instance to drive the forwarding state for multiple data service VPLS instances, reducing the number of control instances required. In addition, processing and responding to PVST and PVST+ packets is supported.

For more information, see the 7210 SAS-D,E,K Services user guide.

Support for LAG on access ports and access-uplink ports

Platforms Supported: 7210 SAS-K

In prior releases, 7210 SAS-K supported LAG on access-uplink ports in active/active mode. With this release, support has been enhanced to include support for active/standby mode to provide 1+1 protection using LAG.

In addition, LAG support has been extended to include access ports providing an option to support redundant Ethernet connections at the customer handoff point for improved service resiliency and availability.

For more information, please refer the 7210 SAS-D,E,K Interfaces user guide.

SUPPORT FOR DOT1X (802.1X)

Platforms Supported: 7210 SAS-K.

This release includes support for 802.1x (Dot1x). Only Dot1x server mode (7210 SAS acts as the authenticator) providing an option to implement port authentication towards the customer handoff points.

For more information, please refer the 7210 SAS-D,E,K Interfaces user guide.

SUPPORT FOR FRAME-BASED ACCOUNTING (FBA)

Platforms Supported: 7210 SAS-K.

This release supports Frame Based Accounting (FBA) which allows the shapers in the device to account for the Ethernet layer overhead (i.e. IFG (12 bytes) + Premable (8 bytes) = 20 bytes) while distributing the available bandwidth to the service queues. FBA can be enabled or disabled per port. When enabled on a port, it affects all the queue shapers, both ingress and egress, configured on that port.

For more information, please refer to the 7210 SAS-D,E,K QoS user guide.

SUPPORT FOR IP AND MAC CLASSIFICATION FOR SAP INGRESS QOS

Platforms Supported: 7210 SAS-K

In prior releases, only Dot1p and IP DSCP classification was supported for access SAP ingress classification. This release adds support for IP (both IPv4 and IPv6) criteria and MAC criteria for access SAP ingress classification providing more flexibility for identifying traffic flows that need differentiated QoS treatment.

NOTE: Before associating a SAP ingress policy with either IP or MAC criteria, resources must be allocated from the ingress-internal-tcam resource pool. This can be done using the CLI command configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource. The ingress-internal-tcam resource pool is shared by multiple features and resources might need to be taken away from other features before allocating it for use by SAP ingress IP-MAC classification.

For more information about this feature, refer to the 7210 SAS-D,E,K QoS user guide and 7210 SAS-D,E,K Basic systems user guide.

SUPPORT FOR G8032 WITHOUT CCMS

Platforms Supported: 7210 SAS-K

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 and can be used with either access SAPs or access uplink SAPs configured in the VPLS 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.

NOTE: In this release, CCMs cannot be used for fault detection when deploying G8032 based rings. Instead, failure detection is based only on physical loss of link detection.

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

SUPPORT FOR IP DSCP MARKING ON ACCESS AND ACCESS-UPLINK PORTS

Platforms Supported: 7210 SAS-K.

With this release, an 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. If there are more number of tags in the packet, then the IP DSCP value is not modified.

For more information, refer to the 7210 SAS-D,E,K QoS User Guide.

PTP BC AND PTP HYBRID MODE

Platforms Supported: 7210 SAS-K

In prior releases, support was available for PTP OC slave. It has been enhanced with this release to include support for PTP BC and PTP hybrid mode. IEEE default profile for time and frequency is supported for PTP BC.

For more information, refer to the 7210 SAS-D,E,K Basic System Configuration User Guide.

SUPPORT FOR PTP-BASED SYSTEM TIME AND OAM TIMESTAMPS (BETA ONLY)

Platforms Supported: 7210 SAS-D ETR, 7210 SAS-K, 7210 SAS-M (access-uplink and network mode), 7210 SAS-T (access-uplink and network mode), 7210 SAS-X, 7210 SAS-R6.

This release allows PTP to be the source of time for the system and OAM packet time stamping.

For PTP to be the source of time for the system, user need to configure NTP to use ptp as preferred server using CLI “configure system time ntp server ptp prefer”. PTP is used as an NTP Stratum 0 source into the NTP process within the node. A side effect of this allocation as an NTP Stratum 0 source is that the node will begin to advertise itself as being at NTP Stratum 1 level, which may influence NTP peers and clients to change their selected time source.

Note that, 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.

PTP has the capability to achieve a higher accuracy time recovery than NTP and is recommended when one-way delay measurements are to be made across a network. In addition to controlling system time and OAM time stamping,

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

.

ENHANCEMENTS

RELEASE 8.0R2

The following item 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 and 7210 SAS-R6 OS release 8.0R2.

CBS AND MBS VALUES IN KBYTES AND KBITS FOR METERS

Platforms Supported: 7210 SAS-R6, 7210 SAS-D, 7210 SAS-M (both access-uplink and network mode), 7210 SAS-T (both access-uplink and network mode), 7210 SAS-X, 7210 SAS-Mxp.

With this release, user has an option to specify the value of CBS and MBS parameters for meters in units of Kbytes and Kbits. With this, the CLI command provides consistency with support available on other 7x50 products.

NOTE: Support is included for SAP ingress meters, access-uplink port meters, network ingress (both port and IP interface) meters and meter override commands.

For more information, see 7210 SAS Services Guide and 7210 SAS QoS Guide for supported platforms.

RELEASE 8.0R1

The following item 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 and 7210 SAS-R6 OS release 8.0R1.

L2PT SUPPORT FOR CDP, VTP, DTP, ET. AL..

Platforms Supported: 7210 SAS-X

With this release, L2PT support in a VPLS service has been enhanced to include support for CDP, VTP, DTP, UDLD, and PAGP. This enhancement is available for 7210 SAS-X.

NOTE: With this enhancement the maximum number of MAC FIB entries available for PBB IVPLS service is reduced from previous releases.

For more information, see 7210 SAS-X, R6 Services Guide.

IPV6 SUPPORT FOR OUT-OF-BAND ETHERNET MANAGEMENT INTERFACE

Platforms Supported: 7210 SAS-R6

This release allows for use of IPv6 addressing mechanisms and protocols for management of the node.

For more information, see 7210 SAS-M,T,X,R6,Mxp Router Configuration User Guide and 7210 SAS-M,T,X,R6,Mxp System Management User Guide.

CBS AND MBS VALUES IN KBYTES AND KBITS FOR METERS (BETA ONLY)

Platforms Supported: 7210 SAS-R6, 7210 SAS-D, 7210 SAS-M (both access-uplink and network mode), 7210 SAS-T (both access-uplink and network mode), 7210 SAS-X, 7210 SAS-Mxp.

With this release, user has an option to specify the value of CBS and MBS parameters for meters in units of Kbytes and Kbits. With this, the CLI command provides consistency with support available on other 7x50 products.

NOTE: Support is included for SAP ingress meters, access-uplink port meters, network ingress (both port and IP interface) meters and meter override commands.

For more information, see 7210 SAS Services Guide and 7210 SAS QoS guide for supported platforms.

IPV4 FORWARDING

Platforms Supported: 7210 SAS-K

In prior releases, 7210 SAS-K supported only IP host functionality. With this release 7210 SAS-K support is included for IPv4 forwarding for RVPLS interfaces allowing the node to act as an IP router. The IP routing functionality is useful when it is required to forward IP management traffic to other 7210 SAS-K nodes connected to it.

NOTE: It is recommended to use the IP host and router functionality on 7210 SAS-K, only for management of the node. In other words, it cannot be used to deploy services.

For more information, see 7210 SAS-D,E,K Router Configuration User Guide and 7210 SAS-D,E,K System Management User Guide.

BFD FOR TLDP

Platforms Supported: 7210 SAS-M (network mode), 7210 SAS-T (network mode), 7210 SAS-X, 7210 SAS-R6 and 7210 SAS-Mxp

BFD support has been enhanced to support fault detection for TLDP sessions. The BFD sessions use the system IP address or the loopback address and support a minimum timer of 100ms with CPU based processing of BFD messages.

For more information, see 7210 SAS-D,E,K Router Configuration User Guide and 7210 SAS-D,E,K System Management User Guide.

.

KNOWN LIMITATIONS

The following are specific technical limitations that exist in Release 7.0R8 of 7210 SAS OS. The topics are arranged alphabetically.

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]

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]

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.

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).

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, R6 and Mxp, 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, R6 and Mxp, 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, R6 and Mxp, 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 SAS-M, SAS-X, SAS-D SAS-T 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]

IP

- "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]

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. [73552]
- **F**or 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. [77986]
- In 7210 SAS-R6, lag-hashing is not done for BUM traffic in a SAP-SAP case, all the traffic is sent out of the flood-port, "show lag <id> detail" displays flood port.
- 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.

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]

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 SAS-X and SAS-R6. [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-R6, packet drops are seen during global revert MBB in case of SAP, primary path and FRR path reside on different IMM's. [177662, 177643].

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, unlearnt 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 unlearnt traffic. This issue exists with SAS-R6 IMMv1 only. [175208]
- For 7210 SAS-R6 and SAS-Mxp, 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.

OAM

- In case of 7210 SAS-R6 and SAS-T, 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 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-R6, 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 SAS-R6, SAS-X, SAS-M and 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 "l2up"). [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]

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]

QOS

- 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-R6, 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-R6 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-R6, 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-R6 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 remaking 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 vprn 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]

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]

TIMING

- 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]

SERVICES

- 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 ies 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 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 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 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
- A set of source MAC addresses are learned on a spoke SDP. If this traffic ceases to come onto this spoke SDP, and if the same traffic (the source MACs) is received on another spoke SDP that is in STP discarding state, then the MACs do not age out. [77996]
- 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 SAS-M and SAST, SAS-X , Mxp and 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. Contact an ALU representative for more information. 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.
- 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]

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-R6, 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]

STP

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

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]
- In case of SAS-E, SAS-M, SAS-X, SAS-T, 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.
- In case of 7210 SAS-E, SAS-M, SAS-X, SAS-T, It is recommended to shutdown external CF (cf2) or USB (uf1) using CLI “file shutdown <>” before they are physically removed from the system. If removed without shutdown, following error messages may appear. CONSOLE:PLATFORM:UMASS_BBB_Send_Command Bulk Transfer Failed UTC CRITICAL: LOGGER #2002 Base A:PLATFORM:UNUSUAL_ERROR.
- 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]
- The 7210 SAS-R6 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]

KNOWN ISSUES

The following are specific technical issues and limitations that exist in Release 7.0R8 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.

ACLs

- Under tools dump system-resources, Egress Shared ACL Entries shows half of the resources as free even when they are used by Egress Mac+IPv4 ACL Entries/ Egress IPv6 128 bit ACL Entries/ Egress Mac+IPv6 64 bit ACL Entries. This is not observed 7210 SAS-K. [213265-MI].
- Egress filter counter does not increment for action drop on 7210 SAS-R6 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]
- If dot1q and * sap created on same port but belonging to different services, traffic matching dot1q SAP hits the ingress filter attached to * sap only if no ingress filter attached to dot1q SAP. This issue happens in network mode 7210 SAS-M, X, T, R6 and Mxp. [109037-MA].
- When Q1.Q2 and Q1.* sap are configured on the same port, traffic sent with Q1.Q2 tag hits ingress filter attached to Q1.* sap. This issue happens in network mode 7210 SAS-M, X, T, R6 and Mxp. [109133-MA]
- For 7210 SAS-E, an IP filter applied at SAP ingress to filter IGMP packets does not work.

CLI

- Following are some of the known issues with CLI Rollback on 7210 SAS-R6:
  - 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]

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]

IGMP SNOOPING

- In case of SAS-R6, 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]

IP

- In 8.0R1 release hardware based BFD on 7210 SAS-Mxp is supported as BETA only.

- In case of 7210 SAS-R6, 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-R6, directly connected active routes take considerable amount of time to come up. [171954-MI]
- In 7210 SAS-R6, 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]

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. This happens only if the LAG is associated with an IP interface. [85967-MI]

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]

MPLS

- For 7210-R6, packet drops are observed in VPLS/VPRN service when traffic switched over to MBB path. [178203-MA]
- In case of 7210 SAS-R6 and SAS-Mxp, FRR timings are greater than 50ms when saps are also configured on a hybrid port. [178443-MA].
- In case of 7210 SAS-R6, 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]

MIRROR

- 7210 SAS-Mxp allows configuration 2 mirror services.[204034-MI]
- 7210 SAS-K, SAP ingress mirroring works even when dot1x is enabled on the corresponding port.[213378-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]

OAM

- In 7210 SAS-D and 7210 SAS-T access-uplink mode, Down MEP CFM SAA tests Inbound/Outbound values are not proper. [217031-MI]
- In case of SAS-K, Broadcast/multicast traffic dropped on VPLS service where Mac Swap with loopback is enabled. [202211-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 1-DM test results are not proper. [199552-MI]
- 7210 SAS-K, CFM CCM packets not processed when the received CCM packet does not have End TLV. [195613-MI]
- Y.1731 version-1 DM request packets received on Down MEPs are not processed in case of SAS-M. This works on Up MEPs. [203936-M]
- Y.1731 version-1 DM request packets received on Down MEPs are not processed in case of EPIPE service on SAS-T access-uplink mode. [203936-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]

QoS

- 7210 SAS-K “tools dump system-resources” Ingress/Egress Queue count may not display correct count. [218879-MI]
- DSCP mask in sap-ingress qos is not supported though CLI is available to configure.

- Copy of sap-ingress qos policy with ip-criteria 'any' throws error when match entry has dscp string. [215920-MI]
- When network qos port policy configured with unicast and multicast meter for any FC, unicast IP traffic received on port is rate limited by multicast meter in case of 7210 SAS-M and SAS-X. [215777-MI]
- In 8.0R1 release SAP Egress aggregate meters can be configured on up to 128 SAPS. Configuring more than 128 saps can result in system messages and functionality may not work properly. [212188-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]
- 7210 SAS-K, in sap-ingress qos policy match frame-type 802dot3 is not supported though CLI allow. [216111-MI]
- In 7210 SAS-R6 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-R6, 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 unlearnt 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 schedular 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]

TIMING

- 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]
- For 7210 SAS-Mxp combo-ports, when one port is selected as SyncE reference and the port goes operationally down, provisioning other combo port as either reference 1 or reference 2 shows the reference in LOS even if the other port is operationally up.[208776 - MI]
- 7210 SAS-R6, 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-R6 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

SERVICES

- In 7210 SAS-R6, 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-R6, 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 statel 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]

STATISTICS/AC COUNTING

- In case of SAS-R6 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]

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]

SYSTEM

- 7210 SAS-Mxp, SAS-T, 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 AutoInit 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.

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.