

# 7210-SAS M OS 2.0\_R3 SOFTWARE RELEASE NOTES

---

These release notes are for Release 2.0\_R3 of the 7210 SAS-M OS software for 7210 SAS M devices.

---

## RELEASE NOTE ORGANIZATION

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

- [Standard Software Upgrade Procedure](#) on page 2
- [Resolved Issues](#) on page 4
- [New Features](#) on page 8
- [Enhancements](#) on page 11
- [Known Limitations](#) on page 13
- [Known Issues](#) on page 16
- [Supported Hardware](#) on page 25



## SOFTWARE UPGRADE PROCEDURES

The following sections contain information for upgrading to the R2.0R2 software version. In particular, there are sections that describe the following:

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

Note that when upgrading from Release 1.1, ensure that the 7210 SAS M is not power cycled until the following upgrade procedure is completed.

### 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 and configuration files, should reverting to the old version of the software be required.

#### Step 2. Copy 7210 SAS M images to cf1:

The 7210 SAS M image files must be copied to the cf1: device on the 7210 SAS M 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:\2.0R2. 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.

#### 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.

#### 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 from to 2.0R2, a firmware upgrade is required. To ensure the firmware upgrade, execute the **admin reboot upgrade** command.

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

**Step 6.** Upgrade the Golden BOOT Loader (only if all of the above steps were successful).

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

---

## UPGRADING FROM PRIOR RELEASES

**R1.1R7** 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 [Standard Software Upgrade Procedure](#) on page 2.

**R1.1R6** 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 [Standard Software Upgrade Procedure](#) on page 2.

## RESOLVED ISSUES

### RESOLVED IN R2.0R2

The following are specific technical issues that have been resolved in Release 2.0\_R2 of 7210-SAS M OS.

- In a highly scaled configuration, under high CPU load conditions, at times a network LAG flap could cause other LAGs on the 7210 to flap. This issue has been fixed. [96776]
- In a highly IGMP snooping configuration, under high CPU load conditions, at times SDP keepalive failures occurred on certain events (such as a network lag flap). This issue has been fixed. [98042]

### CES

- Online insertion and removal of T1/E1 MDA is now supported.
- Using SNMP CES MDA reboot is now supported.
- The **show port** CLI command output showed DS1/E1 channel status up even when the port was in admin down state. This has been fixed. [93467]
- DS0 channel-group oper status was brought down when CEM SAP was deleted. Adding back the CEM SAP would bring up the channel-group. This has been fixed [97052]

### RESOLVED IN R1.1R8

The following are specific technical issues that have been resolved in Release 1.1\_R8 of 7210-SAS M OS.

- Creating a 2548th filter template along with the action criteria failed. This has been fixed. [95978]
- Changing the meter mode from srctm to trctm in a SAP ingress policy did not reset the PIR value to the default. Now, the PIR value is reset to max. [96267]
- The tools dump system-resources now displays the correct usage of resources for user ACLs. [91662]
- When a copper SFP is removed from a port configured in network mode, the MTU value for the port is changed to 1518. Now, the MTU value is retained. [96998]

### RESOLVED IN R1.1R7

The following are specific technical issues that have been resolved in Release 1.1\_R7 of 7210-SAS M OS.

- When an egress ip-filter with a **drop** default-action is attached to an IP interface, it affected outgoing MPLS traffic through the spoke SDP associated with that IP interface. This issue was seen in 1.1R6 and is fixed. [92819]
- When an egress IP filter is attached to an IP interface, it affected incoming IP traffic through that IP interface. This has been fixed. [88816/92295]
- Pinging with the **do-not-fragment** option for certain packet sizes failed. This has been fixed. [92977]
- Spurious traps **No NTP server available**/'NTP server available' were being sent out at times. This has been fixed. [93106]

- Ingress packet counts for 10Gig ports stopped incrementing under following events:
  - Shutdown/no shutdown of a primary LSP over the 10Gig ports.
  - Deconfiguring and configuring ports under IP interfaces configured on 10Gig LAG.
 This is fixed and counters are updated post the events. [91643]
- The **show router isis database id detail** command did not show the correct bandwidth usage for different priorities. This has been fixed. [91163]
- Occasionally, on 10Gig ports, the following console messages appeared when line rate or close to line rate traffic is sent.
 

```

"CRITICAL: LOGGER #2002 Base 1:MDADRV:UNUSUAL_ERROR
"xmacPollStats: MDA 1, port 25, stat 6 bogus read
2b4f72d0000004cb, last read was 5d439a1c000004ca, s/w count
dcba5d97000004ca"
      
```

 [89089/91859]. This issue has been fixed.
- IGMP packets arriving on a default SAP belonging to a VPLS service with IGMP snooping enabled were being dropped. Now they are forwarded. [87029]
- MPLS packets with three labels were remarked to the default EXP while egressing. This issue is fixed. [81956]
- Single tagged EFM OAM packets were being processed by the CPU. Now they are forwarded in hardware. [82634]
- RADIUS packets went out with the system IP address as the source IP address even though a source address was configured for the RADIUS application. Now, they go out with the configured source address. [92018]
- When IGMP packets received from peers were forwarded to another peer, an extra 4 bytes was added in the Layer 2 frame (the IGMP packets were successfully processed though). Now, the extra 4 bytes are not added. [91891]
- SAP egress stats did not increment when an egress IP filter was applied on an IP interface. Now they increment. [90327]

## RESOLVED IN R1.1.R6

The following are specific technical issues that have been resolved in Release 1.1.R6 of 7210 SAS-M OS.

### ACLs

- Under management access filters, the **src-port match-criteria cpm** option has been removed as it is not supported. [87556]
- Egress filters applied on a .1q SAP of service (A) would match packets (with the same .1q tag) of a different service (B) destined to a .1q\* SAP on the same physical port. Counters associated with the filters also incremented. This was applicable only when an ingress SAP of service (B) had either a null or .1q\* encapsulation. This issue is fixed. [84135] In order to fix this issue, egress ACLs are restricted to 256 for both **mac** and **ip** match criteria. Refer to [Upgrading From Prior Releases](#) on page 3.

- IP** • Disabling and enabling re-distribution of external routes in OSPF did not re-populate NSSA LSAs. This issue is fixed. [78868]

- QoS** • In a SAP ingress policy, modifying the MAC or IP match criteria did not take affect. Traffic continued to hit the old rule. This is fixed. [91295]
- When EFM was put in local or remote loop-back mode, and then reverted back, remarking of DSCP values for IP traffic was incorrect. This is fixed. [90696]

**STATISTICS/ACCOUNTING**

- When traffic was destined to a .1q\* SAP the egress SAP statistics did not increment if there was another .1q SAP on the same port (belonging to a different service) and the .1q tag of the traffic was matching this .1q SAP. This was applicable only when an ingress SAP had either a null or 1.q\* encapsulation. This is fixed. [84162]

**RESOLVED IN R1.1.R5**

The following are specific technical issues that have been resolved in Release 1.1.R5 of 7210 SAS-M OS.

- QoS** • When the description for **ip-criteria** in the SAP ingress QoS policy context is changed, the command may return an error message indicating that the number of resources (classifiers and/or meters) have been exceeded for that policy. This issue has been resolved. [88267/88379]
- In network ingress, metering and accounting did not take place under an IP interface if the spoke SDP is either configured with the **vc-type vlan** or has **force-vlan-vc-forwarding** enabled. This issue has been resolved. [86910/86844]
- In network ingress, dynamic updates of LSP-EXP to FC mapping did not take affect when the spoke SDP is either configured with **vc-type vlan** or has **force-vlan-vc-forwarding** enabled. This issue has been resolved. [86842/86968]

**SERVICES**

- On a spoke SDP, when **force-vlan-vc** is enabled with a **vlan-vc-tag** of 0 (zero), packets egressing out of the spoke did not contain the VLAN value 0. This has been resolved and the packets egress out with a VLAN value of 0 under such a configuration. [88269]
- The **OAM-Learned Count** field has been introduced in this release in the output of **show service id service-id fdb**. This field allows the user to identify the number of MACs installed in the FDB for that service through the **mac-populate** command. [88489]

**RESOLVED IN R1.1.R4**

The following are specific technical issues that have been resolved in Release 1.1.R4 of 7210 SAS-M OS.

- ACLs**
- If there were a large number of filters attached to a number of SAPs, and if the box is rebooted with such a configuration, filter counters may not increment when the traffic matches relevant entries. This issue is fixed in this release. [81290]
  - Filter MAC entry creation using SNMP failed if the tMacFilterParamsEtherType is included while creating the entry. Now filter creation with **ether-type** value works through SNMP. [86962]
- LAG**
- LACP can now be enabled only on a maximum of 12 ports. [86730]
- SERVICES**
- In a scenario where a VPLS service and an Epipe service are using SAPs on the same port and if the Epipe SAP happens to be a default SAP or an explicit null SAP, a **shutdown** of the default SAP (or the explicit null SAP) or a **shutdown** of the Epipe service would stop traffic flowing through the VPLS service. This issue is fixed. [87275]
- STATISTICS/  
ACCOUNTING**
- Spoke SDP stats of the primary spoke SDP in an Epipe service resets to 0 (zero), when it is made non-primary (another spoke SDP is made primary in the service). This issue is fixed, and the stats do not reset on such an event. [87023]
- SYSTEM**
- The **monitor port** and **monitor lag** commands showed incorrect statistics. Now they show correct statistics. [83757/86347]
- QoS**
- In a SAP ingress QoS policy, matching **frame-type** under **mac-criteria** is not supported. This **frame-type** option has been removed. [75162]
  - Snmpget on tNetworkPolicyType for network QoS policy works now. [87314]
  - The **lsp-exp** value is remarked to 0 (zero) rather than the default for the FC when executing a **no lsp-exp-in/out-profile** command in the egress of a network IP interface QoS policy. This issue is fixed. [86959]
  - DSCP values of customer packets were reset to zero if the packets ingressed through an access SAP (if and only if the box was booted with a startup configuration in which the port is set to access mode). This issue is fixed and the DSCP values are preserved. [87966/87971]
  - Under a network port policy with Dot1p classification, traffic for classifiers with outprofile were momentarily marked in-profile when some other classifier with profile in is changed to out. This issue is fixed. [85482]
  - A copy of a network QoS policy worked only if the source policy had default egress configurations. [87091]. A copy of a network QoS policy with the sub-option overwrite did not work. [87033]. Both of the issues have been fixed.

## NEW FEATURES

### RELEASE 2.0R3

The following items describe features added to 2.0R3 of the 7210 SAS-M OS.

#### SYNCHRONOUS ETHERNET

- The 7210 SAS-M provides network synchronization using Ethernet ports. Synchronous Ethernet is a variant of line timing supported on the Ethernet SFP ports with SFPs that support Synchronous Ethernet. Synchronous Ethernet ports configured for line timing provide the best synchronization performance through a synchronization distribution network. Line timing mode derives the timing information from the Ethernet ports. This mode is immune to any packet delay variation (PDV) occurring on Layer 2 or Layer 3 links. When synchronous Ethernet is enabled on the node, the operator can select an Ethernet port as a candidate timing reference. The recovered timing from this port is then used to time the system. In this release the recovered timing information is used to time only the Ethernet port output.

### RELEASE 2.0R2

The following items describe features added to 2.0R2 of the 7210 SAS-M OS.

#### OUT-OF-BAND ETHERNET MANAGEMENT PORT

- 7210 SAS-M and all its variant platforms supports 10/100 Base-T ethernet management port. This port is available for out-of-band management of the node.

#### LDP

- This feature provides the capability to use Label Distribution Protocol (LDP) for setup of MPLS Label Switched Path (LSP). The MPLS LSPs setup using LDP is available for use by SDP. LDP is a protocol used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish LSPs through a network by mapping network-layer routing information to MPLS LSPs. An LSP is defined by the set of labels from the ingress Label Switching Router (LSR) to the egress LSR. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. A FEC is a collection of common actions associated with a class of packets. When an LSR assigns a label to a FEC, it must let other LSRs in the path know about the label. LDP helps to establish the LSP by providing a set of procedures that LSRs can use to distribute labels.



- Y.1731:**
- Y.1731 is defined by ITU-T. It is built on top of the IEEE 802.1ag, but it has more functionalities compared to the former, such as AIS, performance monitoring, etc. In this release, the following Y.1731 functionalities are supported:
    - ETH-CCM
    - ETH-LB
    - ETH-LT
    - ETH-RDI
    - ETH-AIS
    - ETH-TST
    - ETH-DM including one-way and two-way delay measurements. Software timestamps are used for delay measurement.

**CIRCUIT EMULATION  
(CES) MDA**

- In Release 2.0R1, a new MDA **m4-ds1** with 4 x T1/E1 ports, that supports Circuit Emulation Services (CES) was available only for BETA testing. In Release 2.0R2, it is available for deployment. Brief description of the MDA and CES is available under the heading Release 2.0R1 of this section.

**RELEASE  
2.0R1**

The following items describe features added to 2.0R1 of the 7210 SAS-M OS.

**CIRCUIT EMULATION  
(CES) MDA (ONLY  
FOR BETA TESTING)**

- In Release 2.0R1, a new MDA **m4-ds1** with 4 x T1/E1 ports, that supports Circuit Emulation Services (CES) is available. This MDA is available for use on the 7210 SAS-M 24F, 7210 SAS-M 24F 2XFP and 7210 SAS-M 24F 2XFP ETR platforms. It can be inserted into the expansion slot available on these platforms. CES services is used to emulate TDM services over an IP/MPLS transport network. It allows operators to transport legacy T1/E1 services over an IP/MPLS transport network. This release supports only SAToP services as per RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*, Adaptive Clock Recovery (ACR) is available for clock synchronisation.

**ALARM INPUTS**

- The 7210 SAS-M 24F, 7210 SAS-M 24F 2XFP and platforms supports alarm inputs (dry contact input) that allows the operator to monitor and report changes in the external environment (such as temperature changes, door open/close, etc.). It supports 4 alarm inputs pins.

**TRAFFIC  
ENGINEERING  
GRACEFUL  
SHUTDOWN**

- Provides a method to reroute transit LSPs in a bulk fashion away from the node during maintenance of a node. A PathErr message with a specific error code “Local Maintenance on TE Link required Flag” is sent during maintenance of a node. When an LER receives the message, it performs a make-before-break on the LSP path to move the LSP away from the links/nodes whose IP addresses were indicated in the PathErr message. Affected link/node resources are flagged in the TE database so that other routers will signal LSPs using the affected resources only as a last resort.

**MESH SDP SUPPORT  
IN A VPLS SERVICE**

- Prior releases of 7210 SAS-M supported only spoke SDPs, allowing the node to be deployed only in a H-VPLS topology. With the introduction of mesh SDPs, the node can be deployed in a VPLS topology.

**LACP TUNNELING IN  
A VLL SERVICE**

- The 7210 SAS-M supports tunneling of LACP packets only in a VLL service. LACP tunneling can be enabled for an Ethernet port. When LACP tunneling is enabled, the port will not process any LACP packets but will tunnel them instead. This port cannot be added as a member of a LAG.

**SUPPORT FOR  
PROCESSING PVST  
PACKETS**

- Support for processing PVST packets received on a SAP enabled for STP. It allows the 7210 to interoperate with nodes that use PVST. When a PVST BPDU is received on Dot1q or a QinQ SAP, the packet is processed appropriately to update STP states and PVST BPDUs are sent out of that SAP.

**DIGITAL  
DIAGNOSTICS FOR  
SFPs AND XFPs**

- Digital Diagnostics Monitoring is a capability of some SFPs and XFPs where the transceiver module maintains information about its working status in device registers including:
  - Temperature
  - Supply voltage (for SFPs)
  - TX bias
  - TX output power
  - RX received optical power

The transceiver is also programmed at the factory with Warning and Alarm thresholds for Low and High conditions of the above that can generate system events.

The **show port port-id detail** command has been augmented to display the DDM data at the console, and the DDM data is accessible through SNMP for use with RMON threshold monitoring if traps are desired for custom thresholds.

## ENHANCEMENTS

### RELEASE 2.0R2

The following items describe enhancements added to 2.0 R2 of the 7210 SAS-M OS.

#### SUPPORT FOR MICROSECOND TIMER GRANULARITY FOR OAM AND SAA TOOLS

- In earlier releases, the system could not compute and display packet round-trip times and packet delay for values less than 10millisecond. With this feature enhancement, from the current release, the system will compute and display packet round-trip times and packet delay for values less than 10 millisecond. Most of the OAM tools and SAA tools have been updated to provide display microsecond time intervals.

#### HARDWARE SUPPORT FOR ETH- CFM

- With this release, the 7210 SAS-M uses the inherent capabilities of the hardware for CFM/Y.1731 protocols. Now, CFM/Y.1731 CCM message protocol processing is handled in hardware, allowing use of CCM timers at 100ms time interval for down MEPs.

### RELEASE 2.0R1

The following items describe enhancements added to 2.0 R1 of the 7210 SAS-M OS.

#### IMPROVEMENT IN FRR FAILOVER TIME

- FRR failover time has been improved. With this enhancement FRR failover time is close to sub-50ms. This allows for increasing service availability in MPLS networks using this node.

#### RSTP FORWARDING IN HARDWARE

- RSTP packets with one or more tags received on a null SAP or RSTP packets received with two or more tags on a Dot1q SAP or RSTP packets received with three or more tags on a QinQ SAP are forwarded in hardware. In prior releases, forwarding of such RSTP packets was done in software. Note that RSTP packets with null tag received on a null SAP or RSTP packets received with a single on Dot1q SAP or RSTP packets received with two tags on a QinQ SAP are processed locally, if STP is enabled. QinQ SAPs are supported only in 7210 SAS-E in this release.

#### INCREASE IN SAP AND SERVICE SCALING

- The number of SAPs and services available for use has been improved from the existing numbers. This allows the user to aggregate more number of services per node. The user is provided with an option to allocate QoS resources per SAP, according to service requirements and to balance this with service scaling requirements.

### RELEASE 1.1R8

The following items describe enhancements added to 1.1.R8 of the 7210-SAS M OS.

- ETR PLATFORM Release R1.1R8 introduces a new platform, 7210 SAS-M 24F 2 XFP ETR, that supports an extended temperature range. The new platform is a variant of the

7210 SAS-M 24F. It is available in an ETSI-conforming chassis which is 1.5RU in height and 19-inch in width and supports an extended temperature range of -40°C to +65°C. Its a MPLS-capable switch and supports line-rate switching on all its ports. It supports 24 x 100/1000 SFP ports and 2 x 10G XFP ports and provides an additional expansion slot for future use. It requires the use of extended temperature range power supplies. It is feature compatible with 7210 SAS-M 24F 2XFP.

## RELEASE 1.1R7

The following items describe enhancements added to 1.1 R7 of the 7210 SAS-M OS.

### 10G PLATFORM

This release supports a new platform, the 7210 SAS-M 24F 2XFP. This platform is a variant of the 7210 SAS-M 24F and supports an additional 2 x XFP ports, along with 24 x 100/1000 SFP ports. It is an MPLS switch/router capable of line-rate forwarding. It is feature compatible with the 7210 SAS-M 24F. Refer to the **7210-SAS-M OS Services Guide** for relevant CLI commands.

### AC/DC POWER SUPPLIES

7210 SAS-M and 7210 SAS-M 24F 2XFP platforms can accept and operate with a combination of AC and DC power supplies simultaneously.

### SERVICE MTU DISABLE

This feature provides a mechanism to disable service MTU validation in a service. This is required for interoperability with third-party devices. Service MTU Disable also allows the operator to work around the limitation (see [Known Limitations on page 13](#)) in 7210 SAS-M when **service-mtu-check** is enabled. Refer to the **7210-SAS-M OS Services Guide** for relevant CLI commands.

### LLF

The 7210 SAS-M now supports Link Loss Forwarding in Epipe services. Link Loss forwarding enables fault propagation between the two endpoints of the Epipe service. The two endpoints in the Epipe service can be null SAPs or a null SAP and spoke SDP. Refer to the **7210-SAS-M OS Services Guide** for relevant CLI commands.

### IGMP SNOOPING

- In R1.1R6, 7210 SAS M introduced IGMP snooping in a VPLS service support. 7210 SAS M supports IGMPv1 and IGMPv2. IGMP snooping allows the system to distribute multicast streams to only interested receivers in a Layer 2/bridging domain. Refer to the **7210-SAS M OS Services Guide** for relevant CLI commands.

## KNOWN LIMITATIONS

- The **show eth-cfm mep mep-id domain md-id association ma-id** command would not display Last CCM Rx-Time, Last CCM Seq Number, CCM Tx Count.
- ETH-CFM CCM Sequence Check for out-of-sequence CCMs and Last CCM Sequence Number tracking is not done for received CCMs.
- ETH-CFM CCMs with incorrect DMAC are not dropped by hardware.
- ETH-CFM CCMs transmitted will always have the Sequence Number as zero.
- ETH-CFM CCMs with invalid Data TLV-s , Invalid Port/Interface TLV-s do not cause CCM Errors
- Under a high CPU load, ETH CFM defect reporting and clearing may be delayed.
- On access ports, protocol packets of EFM, LACP, Dot1x, and DWL, are counted as part of SAP statistics if a SAP configured on an access port of **encap type null, explicit-null** or **dot1q-\***. [95361]
- FRR failover time for unlearned traffic (such as broadcast, unknown-unicast, and multicast traffic types) will not be under sub-50ms.
- In a VPLS service, when IGMP snooping is enabled, the multicast replication is based on Layer 2 MAC addresses.
- IGMPv3 is not supported.
- In VPLS and VLL services, port MTU checks are performed only at the ingress. [92910]
- 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).
- 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 v2 leaves are sent upstream when V1 reports expire. [88271]
- IGMP reports and leaves received from multicast clients/receivers are sent out of the static Mrouter port configured in the service even when the service does not have a dynamic querier port.
- DC source (-48V) failure cannot be detected if both sources (-48VA and -48VB) are connected to the 7210 SAS M.
- When packets are received on a Dot1q SAP, the incoming packet length (frame length along with the length of the encapsulation tag) is used to match against the service MTU configured. If the incoming packet length is greater than the service MTU configured, the packets are dropped. Therefore, while configuring the service MTU, one should account for the SAP encapsulation and spoke SDP VC-VLAN tags. [79679]
- Dual rate SFPs (3HE04116AA and 3HE04117AA) connected to GigE SFP require autonegotiation to be disabled to operate in 1G mode. [78737]

- SAP ingress meters (counters) are incremented for packets dropped by a filter on that SAP. [70878]
- 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.
- In an Epipe service, traffic is not switched if the source MAC is a multicast or broadcast address [71437]
- The system accepts packets with sizes exceeding the port MTU by 4 bytes if the packet is VLAN tagged. [75221]
- Tagged BPDUs on a null SAP are forwarded by the software. [71656]
- A MAC will not age out as long as STP BPDUs from that source are received, although data traffic is not present. [71658]
- 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 in-band RADIUS servers, the source address (IP address) is used as the NAS IP address, otherwise the IP address of the system interface is used.
- **default.cfg** is a file name reserved by the system. Do not create a file with this name in the root directory. [76972]

- For a LAG configuration with more than one port, if a meter configuration does not specify a CBS value, some packets may be marked yellow and be treated accordingly when buffer management begins.

The solution is to increase the CBS setting whenever the CIR configured for the meter is greater than 1Gbps. For a LAG with two ports, a CBS value of 64Kbits is recommended. [72497]

- 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]
- 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]
- MACs that are already learned do not age out after a filter is added to drop packets from those MACs. [73370]
- When egress port mirroring is enabled, packets that are dropped due to a split horizon check are mirrored.
- Packets that are dropped due to egress queue drops are mirrored when port egress mirroring is enabled.
- Packets that are dropped due to egress filters are mirrored when port egress mirroring is enabled.
- 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]
- 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]
- 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]
- Ethernet ports that use copper SFPs cannot participate in a Synchronous Ethernet Network.

## KNOWN ISSUES

The following are specific technical issues and limitations that exist in Release 2.0\_R3 of 7210 SAS-M OS. The topics are arranged alphabetically.

- ACLs**
- An IP filter applied at SAP ingress to filter IGMP packets does not work. [89256]
  - A **mgmt-access-filter** with a **deny-host-unreachable** action sends “Destination Net Unreachable” instead. [73676]
  - 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]
- CES**
- CES packets, when sent over a LAG, will get load distributed causing packet re-ordering. In this release, CES is not supported in a configuration where the packets egress out of a LAG. [97616]
  - 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]
  - SAP ingress statistics in a CPIPE service does not increment when LOS is reported on a DS1/E1 port. [98828]
  - 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]
  - TDM ports cannot participate in a split horizon group (although it is user configurable). It is not a supported feature. [101695]
- CLI**
- 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]
  - In the Telnet session, the output for **tools dump** is not properly displayed. [80661]
- IP**
- The CLI commands under the context **config>router** for the **management** routing-instance are not supported. [101636]
  - Sometimes **Packet failed authFailure authentication** messages are seen on OSPF MD5-enabled interfaces. They are harmless. [90099]
  - DNS packets do not use the source IP configured for DNS applications. [75420]
  - IP packets with options are not forwarded.



- 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]
- Performing an SNMP GET on an interface description may not return the correct description of that interface if the same CLI command was used to create the interface and its description (for example, **configure router interface** *ifName* **description** *ifDesc*). This is not an issue if the interface was created first, followed by the description. [65622]
- The **show router interface** command output incorrectly truncates the **Port/SapId** field for any interface when the string length of this field is more than 14 characters. [66751]
- 7210 SAS M does not support secondary IP interfaces. However, these are configurable through SNMP, and should be avoided. [76848].
- 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]
- SNMP query of the vRtrActiveArpEntries object does not return the correct value. The CLI reports the correct number of ARP entries. [80788]

**LAG**

- When all of the traffic is egressing out of one port of a two-port LAG, the traffic switchover time onto the other port (due to a link down event) is higher compared to when traffic is shared between the two ports. [86060]
- 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]
- IP-based traffic is not load balanced between the ports of a LAG with the traffic flows from spoke SDP-to-spoke SDP or spoke SDP-to-SAP. [84791]

**MANAGEMENT**

- 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]
- 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]
- LACP flaps when starting an SSH server (no shutdown of SSH server). [75648]
- SNMP operations on some unsupported SNMP MIBs might succeed.
- There is currently no **show** command to display the current values of password hash settings. [32747]
- When logged in with the Telnet client supplied with Windows and issuing an SSH command to a remote server, an unexpected carriage return may be sent causing the SSH connection to fail. [50657]
- 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]

- 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]
- SNMP walk of vRtrConfEntry shows the vpls-management and management instance as active even though these are not currently supported. [76832]
- SNMP query of the following operational rates does not return the correct values. The value returned is 0. CLI reports correct operational values. [76853].
  - tAccessEgressQueueOperPIR
  - tAccessEgressQueueOperCIR.
- 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]
- The 7210 SAS M does not support storing more than 500 events for log destination's memory and SNMP. Although the CLI and MIB allows up to 1024 to be configured, it is recommended not to exceed 500.

#### MIRROR

- Up to 25 SAPs can be ingress mirrored. [98026]
- Log events are not generated for mirror application. [72100]
- CPU-generated STP packets are not mirrored when port egress mirroring is enabled. [84200]
- Ingress QoS policies applied for forwarded traffic will also be reflected on its mirrored traffic. [73951]

#### MPLS

- 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]
- In certain conditions, VPLS traffic may be affected due to a network LAG flap. This was seen only with LDP tunnels. [101142]
- The operational state of LDP will go down with the reason code "IOM Failure" if the number of /32 prefixes learnt go beyond 255. [98085]
- The **show router rsvp interface interface-name detail** command displays incorrect Auth Rx Seq Num and Auth Tx Seq Num values. [86903]

#### OAM

- The **show eth-cfm mep mep-id domain md-id association ma-id** command would not display CCM ERROR, CCM XCON frames in its output.
- The **show eth-cfm mep mep-id domain md-id association ma-id remote-mep rmep-id** command would not display some TLVs and details.
- ETH-CFM session on standby pseudowire does not come up in a VLL pseudowire redundancy configuration. [100594]
- When MEPs are configured on active and standby spokes in a VLL pseudowire redundancy scenario, on a failover, the MEPs on the new active spoke will have the defect flag set. The issue is not seen if the MEP is configured only on the active spoke. [100816]
- ETH-CFM CCM/LTM priority is not correctly mapped to the EXP values when egressing out of SDPs.

- ETH-CFM RDI will not be generated for defMacStatus [98332]
- ETH-CFM defect is not reported on booting the 7210 with a configuration that has the ETH CFM session over a LAG with no member ports [100813]
- In certain conditions, while creating an MEP in a service and a domain, the following messages may appear:
  - IOM:attach\_meter\_and\_action Entry exists.
  - IOM:program\_default\_ifp\_entry Attach Meter failed for EID ... [100784]
- Lsp-ping/trace to ldp prefix with path-destination any other than 127.0.0.1 would fail. [99489]
- The **probe-count** option with value greater than 1 used in **oam lsp-trace** will not send more than one echo-request. It works appropriately with SAA. [85119]
- ICMP pings with higher packet sizes sent at higher rates will fail. [77611]
- Under a heavy load condition, CFM/EFM/LACP may flap due to any of the following triggers [74223]:
  - 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.
- The minimum dead-interval time supported for an EFM OAM session is 2 seconds.
- OAM DNS lookups do not work correctly unless the full DNS name is provided. [54239, 54689]
- Efmooam sessions flap under the following conditions:
  - 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]
- 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]
- 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]
- 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]
- For SAA tests for **mac-trace**, the result shows success for unknown MAC addresses. [81891]
- The timestamps are all 0s when **cpe-ping** is performed with SAA. [81726]

#### QoS

- Packets received on a MPLS LSP setup using LDP, EXP to FC classification along with profile identification for color aware policing on network ingress, will use the default network policy (type - ip-interface). Though the ingress classification (and also profile) is not user configurable, network ingress meters are available for user configuration. Users

can configure meters per network IP interface using the existing network policy. In other words, for packets received on MPLS LSPs setup using LDP, the configured parameters for FC classification is ignored and only meter configuration is used from the network policy (type ip-interface) [99401]

- A maximum of 15 network QoS policies of type **ip-interface** with unique mapping of FC to EXP values can be created and these policies must be shared among 32 IP interfaces. When the IP interface is down, attaching or removing a different QoS policy to the IP interface multiple times results in exhaust of QoS resources and new QoS policy cannot be attached. [93295/101423/97873]
- Operational values for some user configured values of ERL rate are incorrect in the **show pool** output of a port. [90549]
- L2PT tunneled STP packets do not have the appropriate MPLS EXP bits set. They are set to zero. [90580]
- SRTCM meter configured on 10Gig ports with the CBS value default and MBS value of 33547 or greater, the ingress metering is not correct. [89089]
- When a 10Gig 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. [89963, 89094, 89095]
- 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]
- 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]  
The workaround: In the case of a NULL SAP, remarking can be disabled to preserve the Dot1p bits.
- 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]
- 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]
- The keyword **max** in the meter command configuration does not set the rate properly. The semantics of this keyword should set the metered rate to the maximum rate applicable to the interface over which the SAP attached to the meter is created. However, **max** sets the rate to 4Gbps. Use the applicable numeric value instead. [75845].
- 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]
- An access egress policy with remarking enabled in VPLS services with the **vc-type** equal to **vlan** or **force-vlan-vc-forwarding** does not remark Dot1p values. This is applicable to all traffic ingressing a spoke SDP and egressing out a Dot1q SAP. [81761]
- Dynamic updation of **tos-marking-state** at the ingress ip interface of an LSR (that uses RSVP tunnels) does not take effect. [102733]

- 
- RADIUS**
- In defining RADIUS Vendor Specific Attributes (VSAs), the **TiMetra-Default-Action** parameter is required even if the TiMetra-Cmd VSA is not used. [13449]
- SERVICES**
- Ingress SDP statistics are accounted against the primary spoke SDP, even if traffic is received on other secondary spoke SDP. (93627).
  - In a scaled setup with spoke SDPs configured as mrouter port multicast traffic for some services is not forwarded over some spokes when the LAG flaps. Remove and re-configure the mrouter port for the spoke SDP restores the traffic. [97851]
  - Disabling IGMP snooping in a service causes static entries on a spoke SDP to be removed. [91105]
  - 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]
  - 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 Mrouter port. The work-around is to remove the static Mrouter configuration and add it back; or execute a **shutdown** command and a **no shutdown** command of the spoke SDPs. [91856]
  - 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]
  - Packets larger than port MTU are learned and are accounted for by rate limiters. However, they are dropped as expected. [73497]
  - 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]
  - 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]
  - 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]
  - Changing the **fdb-table-size** on multiple services simultaneously (for example, by a script) may affect MAC learning in some services. [82362]
  - 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]
  - Changing the **fdb-table-high-wmark** and **fdb-table-low-wmark** values more than once, will stop generating any notification. The workaround is to delete and recreate the service. [81188]
- STATISTICS/  
ACCOUNTING**
- LACP and CFM protocol packets are not shown in the output packets column of the **show lag lag-id statistics** command. [77986]
-

- STP**
- RSTP convergence fails if **force-vlan-vc** is enabled on a mesh SDP. [94928]
  - 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]
  - If an ingress CPU forwarded RSTP/L2PT/PVST packet rate exceeds 128kbps on access, 265kbps on network, RSTP flaps. [73189]
  - 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]
- SYSTEM**
- 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]
  - 10Gig alarms, “no-frame-lock” and “high-ber”, are not supported on the 7210 although they are configurable.
- SYNCE**
- 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.
  - Reference switch is based on LOS and not based on signal degradation.
- SYSTEM**
- 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]
  - 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]
  - 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]
    - Addition and deletion of port (or LAG port) to ip-interface. [84165/86007]
    - Enabling an SSH server. [84166] 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.
  - 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]
  - Traffic drops are seen in line rate bi-directional conditions where BUM traffic contributes to greater than or equal to 50% of line rate. [85380]

- 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]
- 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]
- 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]
- When the password aging option is enabled, the reference time is the time of the last boot, not the current time. The password expiry will also be reset on every reboot. [64581]
- 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]
- 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].
- 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].
- Inserting and removing SFPs in rapid succession cause the following message to display on the CLI session.

**SFP/XFP Checksums do not match.**

If this message is seen, execute a **shutdown** command followed by a **no shutdown** command on the offending port to resolve the issue. [76935].

- 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].

- During an auto-boot through DHCP when a ping-address is used in the BOF, the system may enter an inconsistent state during the ping-check phase. This was observed only once during which the user was viewing certain debug commands. [76858]
- When 1G fiber SFPs are configured with a 100Mbps speed may result in un-predictable behavior. Do not set the speed to 100Mbps for 1G fiber SFPs. [81948]
- For dual rate fiber SFPs to operate successfully, both sides should be configured with the correct speeds. [82194]
- When using 100FX SFPs, autonegotiation should be disabled. [99132]
- The number of files in the root directory is limited to 100. As a possible workaround, create a directory in the root directory to use to save/store files. [75227]

- HARDWARE**
- 10Gig XFPs with the following part codes are not functional on the 7210. It is recommended not to insert them on the 7210. [99074]
    - SXP3102EX-W
    - SXP3100SX-A8
    - SXP3102ZX-A8
    - AFBR-720XPDZ
    - SXP3101LX-A4
  - The USB interface is not supported on the 7210 SAS M.
  - Only one of the two 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.



## SUPPORTED HARDWARE

The following section lists supported SFPs..

| Supported SFPs     |   |
|--------------------|---|
| 100FX              | 3HE00869CA 1-port 100FX BiDi-U Tx 1310nm Rx 1490nm 10km 1310nm<br>3HE00869CB 1-port 100FX BiDi-D Tx 1490nm Rx 1310nm 10km 1310nm<br>3HE04323AA 1-port 100BASE BiDi-U TX:1310 nm, RX: 1550 nm, 40km<br>3HE04323AB 1-port 100BASE BiDi-D TX:1550 nm, RX: 1310 nm, 40km<br>3HE00266AA 1-port 100fx 25km 1310nm<br>3HE00024CA 1-port 100fx-sx 1310nm<br>3HE04524AA 1-port 100BASE-FX 10km, 1310 nm<br>3HE01454CA 1-port 100BASE-FX 40 km, 1310 nm   |
| GigE               | 3HE00029CA 1-port 1000base-Lx 1550nm 70km<br>3HE00070CA 1-port 1000BASE CWDM 1471 nm, 120Km<br>3HE00070CB 1-port 1000BASE CWDM 1491 nm, 120Km<br>3HE00070CC 1-port 1000BASE CWDM 1511 nm, 120Km<br>3HE00070CD 1-port 1000BASE CWDM 1531 nm, 120Km<br>3HE00070CE 1-port 1000BASE CWDM 1551 nm, 120Km<br>3HE00070CF 1-port 1000BASE CWDM 1571 nm, 120Km<br>3HE00070CG 1-port 1000BASE CWDM 1591 nm, 120Km<br>3HE00070CH 1-port 1000BASE CWDM 1611 nm, 120Km<br>3HE00867CA 1-port 1000base-ex 1310nm 40km<br>3HE01389CA 1-port 1000base-ezx 1550nm 120km<br>3HE00027CA 1-port 1000base-sx 850nm<br>3HE00028CA 1-port 1000base-Lx 1310nm 10km<br>3HE00062AA 1-port 1000base-TX<br>3HE00868CA 1-port 1000base-BX-U BiDi Tx 1310nm Rx 1490nm 10km<br>3HE00868CB 1-port 1000base-BX-D BiDi Tx 1490nm Rx 1310nm 10km<br>3HE04324AA 1-Port 1000BASE-BX-U BiDi TX: 1310 nm, RX:<br>1490nm,40km<br>3HE04324AB 1-port 1000BASE-BX-D BiDi TX: 1490nm, RX:<br>1310nm,40km<br>3HE04939CA 1-port 1000Base CWDM-1471 40Km<br>3HE04939CB 1-port 1000BASE CWDM-1491 40Km<br>3HE04939CC 1-port 1000BASE CWDM-1511 40Km<br>3HE04939CD 1-port 1000BASE CWDM-1531 40Km<br>3HE04939CE 1-port 1000BASE CWDM-1551 40Km<br>3HE04939CF 1-port 1000BASE CWDM-1571 40Km<br>3HE04939CG 1-port 1000BASE CWDM-1591 40Km<br>3HE04939CH 1-port 1000BASE CWDM-1611 40Km |
| 100/1000 Dual Rate | 3HE04116AA 1-port 100/1000base FX dual rate SGMII 1310nm 2Km<br>3HE04117AA 1-port 100/1000base LX dual rate SGMII 1310nm 10Km   |

|                |  |
|----------------|--|
| <p>10G XFP</p> | <p>3HE00876CA 1-port 10GBASE-ER , 1550 nm, 40km<br/> 3HE04325AA 1-Port 10GBase 1550nm, 120km<br/> 3HE00564CA 1-port 10GBASE-LW/LR 1310 nm, 10km<br/> 3HE00876AA 1-port 10GBASE-ER 1550 nm, 40km<br/> 3HE01545CA 1-port 10GBASE-ZR 1550nm, 80km<br/> 3HE05832CA 3HE00566CA 1-port 10GBASE-SW/SR 850 nm, 26 to 300 meters<br/> 3HE05831CA<br/> 3HE01545CA 1-port 10GBASE-ZR 1550nm, 80km<br/> 3HE00564AA 1-port 10GBASE-LW/LR 1310 nm, 10 km<br/> 3HE00566CA 1-port 10GBASE-SW/SR , 850 nm, 26 to 300 meters,<br/> 3HE02717AA 1-port 10GBASE-DWDM , 1560.61 nm/Channel 21 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AB 1-port 10GBASE-DWDM , 1559.79 nm/Channel 22 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AC 1-port 10GBASE-DWDM , 1558.98 nm/Channel 23 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AD 1-port 10GBASE-DWDM , 1556.55 nm/Channel 26 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AE 1-port 10GBASE-DWDM , 1555.75 nm/Channel 27 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AF 1-port 10GBASE-DWDM , 1554.94 nm/Channel 28 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AG 1-port 10GBASE-DWDM , 1552.52 nm/Channel 31 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AH 1-port 10GBASE-DWDM , 1551.72 nm/Channel 32 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AJ 1-port 10GBASE-DWDM , 1550.92 nm/Channel 33 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AK 1-port 10GBASE-DWDM , 1548.51 nm/Channel 36 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AL 1-port 10GBASE-DWDM , 1547.72 nm/Channel 37 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02717AM 1-port 10GBASE-DWDM , 1546.92 nm/Channel 38 (100GHz ITU Grid, C-Band), 80 km,<br/> 3HE02718AA 1-port 10GBASE-DWDM , 1561.42 nm/Channel 20 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AB 1-port 10GBASE-DWDM , 1557.36 nm/Channel 25 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AC 1-port 10GBASE-DWDM , 1553.33 nm/Channel 30 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AD 1-port 10GBASE-DWDM , 1549.32 nm/Channel 35 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AE 1-port 10GBASE-DWDM , 1543.73 nm/Channel 42 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AF 1-port 10GBASE-DWDM , 1542.94 nm/Channel 43 (100GHz ITU Grid, C-Band), 80 km<br/> 3HE02718AG 1-port 10GBASE-DWDM , 1542.14 nm/Channel 44 (100GHz ITU Grid, C-Band), 80 km</p> |
|----------------|--|

---

|                |   |
|----------------|---|
| 10G XFP (cont) | 3HE02718AH 1-port 10GBASE-DWDM , 1541.35 nm/Channel 45 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AJ 1-port 10GBASE-DWDM , 1539.77 nm/Channel 47 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AK 1-port 10GBASE-DWDM , 1538.98 nm/Channel 48 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AL 1-port 10GBASE-DWDM , 1538.19 nm/Channel 49 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AM 1-port 10GBASE-DWDM , 1537.40 nm/Channel 50 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AN 1-port 10GBASE-DWDM , 1535.82 nm/Channel 52 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AP 1-port 10GBASE-DWDM , 1535.04 nm/Channel 53 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AR 1-port 10GBASE-DWDM , 1534.25 nm/Channel 54 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AT 1-port 10GBASE-DWDM , 1533.47 nm/Channel 55 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AU 1-port 10GBASE-DWDM , 1531.90 nm/Channel 57 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AV 1-port 10GBASE-DWDM , 1531.12 nm/Channel 58 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AW 1-port 10GBASE-DWDM , 1530.33 nm/Channel 59 (100GHz ITU Grid, C-Band), 80 km<br>3HE02718AX 1-port 10GBASE-DWDM , 1529.55 nm/Channel 60 (100GHz ITU Grid, C-Band), 80 km |
|----------------|---|