# 7210-SAS X OS 2.0_R5
# SOFTWARE RELEASE NOTES

These release notes are for Release 2.0_R5 of the 7210 SAS-X OS software for 7210 SAS X devices.

## RELEASE NOTE ORGANIZATION

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

# SOFTWARE UPGRADE PROCEDURES

The following sections contain information for upgrading to the 2.0_R5 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 and configuration files, should reverting to the old version of the software be required.

**Step 2.** Copy 7210 SAS X images to cf1:

The 7210 SAS X image files must be copied to the cf1: device on the 7210 SAS X 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:\2.0R5. 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 to 2.0R5, execute the **admin reboot upgrade** command. Note that, executing 'admin reboot upgrade' command system will upgrade the bootrom if required.

**Step 6.**

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

**Step 7.** Upgrade the Golden BOOT Loader (only if all of the above steps were successful).

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

# KNOWN LIMITATIONS

- 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.
- 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).
- 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]
- 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]
- Only egress rate command can be used on mirror destination. sap-egress qos policy cannot be used with mirror destination (null sap).

- 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_R5 of 7210 SAS-X OS. The topics are arranged alphabetically.

**ACLs**

- SAP Egress Mac filter counters doesnot increment for packets with MPLS payload, filter functionality works fine. [107279]
- Maximum of 1024 IP Filter and 1024 MAC Filter entries are supported, but CLI allows creation of 8k MAC entries. [105669] [105900]
- An IP filter applied at SAP ingress to filter IGMP packets does not work. [89256]
- 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]

**CLI**

- The execution of configuration file saved with CLI 'admin save detail' errors out.
- 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**

- Disabling OSPF under Management Routing Instance brings down the OSPF Neighbors in Base Routing instance. [106441]
- 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]
- 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 X 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**

- Sometimes, it is observed that, in lag configured with subgroup, non-unicast traffic flows on the standby port after certain events like lag port flaps and lag shut/no shut events. Workaround is to remove subgroup and add it back. [107278]
- 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**

- When SAS-X is the penultimate-hop to an egress LER that advertized an implicit null label, the SAS-X, on processing an incoming single-labeled (IP) packet, sends out the packet towards the egress LER by stripping the label, while maintaining the ether-type as 0x8847 instead of 0x0800. [104649]
- 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 **show router rsvp interface** *interface-name* **detail** command displays incorrect Auth Rx Seq Num and Auth Tx Seq Num values. [86903]

**OAM**

- Software based timestamps are used for Y.1731 delay measurements.
- With implicit null label advertised at the LER and SAS-X as the penultimate hop, lsp-trace ping does not work if network interface is dot1q. [105658]
- The **probe-count** option with value greater than 1 used in **oam lsp-trace and vccv-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]
- Efmoam 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]
- 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**

- When traffic sent from SAP to SAP at line rate packet drops are seen. 64 Bytes frame at 97% or 1518 Bytes frame at 99% go through without any drops. Similar thing is observed with Egress-port rate configuration.[103718]
- Egress rate calculation is frame based. When traffic egressing out of a queue, the configured queue rate accounts for IFG of the frame.[103790]
- When Traffic egressing out of network port, traffic is rate-shaped without taking MPLS 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.[101713]
- When traffic is send 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 also with egress-rate configured on the port [102176]
- During configuration oversubscription of queue CIR results in undesirable behaviour of shaper and schedular.[101715]
- Queue CIR/PIR minimum value can be set to 26, configuring lesser value results in traffic drop.[102578]
- 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]
- When traffic is bursty in nature the default CBS and MBS would not be enough for a burst and there could be packet drops. Hence user may need to increase the CBS and MBS for the queue.[102505]
- Egress rate command used on 10G results in random/unfair traffic distribution Cir is not being met and PIR distribution is random [106307]

- 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.
- Network QOS policy's Copy overwrite command fails when there is a difference in the egress part of the policy. [106874]
- Creation of remark policy using mib ‘tSASRemarkPolicyRowStatus’ without type mib "tSASRemarkPolicyType' results in creating remark policy with non-default type on SAS-X.[106642]
- Service STP BPDU's egressing out of network port are remarked to MPLS-EXP value of Zero when re-marking enabled on network qos policy applied on IP Interface [107033]
- CPU generated traffic is not counted in egress queue counters statistics. [105180]
- If the SAP is attached with non default sap-egress qos policy, before deleting SAP the attached sap-egress qos policy should be removed. Otherwise in some case buffer leakage can happen and 'show pools' shows wrong value for CBS.[106169]
- For multipoint queues If CIR of any queue is configured more than 800Mbps then the PIR serving of other queues is not proper. Workaround is which ever queue exceeding 800 Mbps, for that queue by setting CIR = PIR then the PIR serving is working fine.[107065]
- 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 31network 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. [106718]
- 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]
- 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]
- 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]

- 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]

**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**

- IGMP Snooping in vpls services is not supported. IGMP snooping CLI's should not be turned on.
- BUM traffic sent on a sap are counted in egress queue counts of source sap. [106158]
- 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]
- 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]

**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-Xs 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**

- System deriving clock from 10 gig port with XFP (3HE00566CAAA01) doesn't move to holdover state when the master Dut is admin rebooted, power cycle of master Dut it works fine.[105269]
- SyncE not supported on the following SFP's.
  - Dual Rate SFP's: 3HE04116AA and 3HE04117AA
  - Fast Ethernet SFP's: 3HE00869AB, 3HE01454AAAA02 and 3HE00024AAAA02
- 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**

- Inserting or removal of CuSFP or Dual rate SFP on port 9 or 21, results in adjacent ports flap.[107210]
- 10Gig XFPs with the part code SXP3100SX-A8 is not functional on the 7210. [104832]
- 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. [106397]
- 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**

- The USB interface is not supported on the 7210 SAS X.
- 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 | 3HE01454AA 1-port 100BASE-FX 40 km, 1310 nm<br>3HE04524AA 1-port 100BASE-FX 10km, 1310 nm<br>3HE00024AA 1-port 100FX-SX 1310 nm<br>3HE00869AA 1-port 100BASE BiDi-U TX:1310 nm, RX: 1490 nm, 2-10 km<br>3HE00869AB 1-port 100BASE BiDi-D TX:1490 nm, RX: 1310 nm, 2-10 km |
| GigE | 3HE00867CA 1-port 1000base-ex 1310nm 40km<br>3HE00868AA 1-port 100BASE BiDi-D TX:1550 nm, RX: 1310 nm, 2-10 km<br>3HE00868AB 1-port 100BASE BiDi-D TX:1550 nm, RX: 1310 nm, 2-10 km,<br>3HE00028AA 1-port 1000BASE-LX 1310 nm, 10 km<br>3HE00027AA 1-port 1000BASE-SX 850 nm<br>3HE01389AA 1-port 1000BASE-EZX 1550nm, 120 km<br>3HE00029AA 1-port 1000BASE-ZX 1550 nm, 70 km<br>3HE00070BA 1-port 1000BASE CWDM (120 km) 1470 nm<br>3HE00070BD<br>3HE00070BE 1-port 1000BASE CWDM (120 km) 1550 nm<br>3HE00070BG 1-port 1000BASE CWDM (120 km) 1590 nm<br>3HE00070BH 1-port 1000BASE CWDM (120 km) 1610 nm |
| 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 |
| 10G XFP | 3HE00876AA 1-port 10GBASE-ER 1550 nm, 40km<br>3HE00566CA 1-port 10GBASE-SW/SR 850 nm, 26 to 300 meters<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>3HE00876AA<br>3HE00564AA |