Release Notes

OmniSwitch 10K, 6900

Release 7.3.2.R01

These release notes accompany release 7.3.2.R01 software which is supported on the OmniSwitch OS10K and OmniSwitch 6900 platforms. These release notes provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

ATTENTION: Please refer to the 7.3.2.R01 Prerequisite section for important release specific information prior to upgrading including specific build information for hardware support.

Contents

Related Documentation

These release notes should be used in conjunction with OmniSwitch AOS Release 7 User Guides. The following are the titles and descriptions of the user manuals that apply to this release. User manuals can be downloaded at: http://enterprise.alcatel-lucent.com/?dept=UserGuides&page=Portal

OmniSwitch 10K Series Getting Started Guide
Describes the hardware and software procedures for getting an OmniSwitch up and running.

OmniSwitch 10K Series Hardware User Guide
Complete technical specifications and procedures for all OmniSwitch Series chassis, power supplies, and fans.

OmniSwitch 6900 Series Getting Started Guide
Describes the hardware and software procedures for getting an OmniSwitch up and running.

OmniSwitch 6900 Series Hardware User Guide
Complete technical specifications and procedures for all OmniSwitch Series chassis, power supplies, and fans

OmniSwitch AOS Release 7 CLI Reference Guide
Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.

OmniSwitch AOS Release 7 Network Configuration Guide
Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethe rnet and VLAN configuration), Layer 3 information (routing protocols), security options (Authen¬ticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.

OmniSwitch AOS Release 7 Switch Management Guide
Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).

OmniSwitch AOS Release 7 Advanced Routing Configuration Guide
Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, and OSPFv3.

OmniSwitch AOS Release 7 Data Center Switching Guide
Includes and introduction to the OmniSwitch data center switching architecture as well as network configuration procedures and descriptive information on all the software features and protocols that support this architecture. Chapters cover Shortest Path Bridging MAC (SPBM), Data Center Bridging (DCB) protocols, Virtual Network Profile (vNP), and the Edge Virtual Bridging (EVB) protocol.

OmniSwitch AOS Release 7 Transceivers Guide
Includes SFP, SFP+, and QSFP transceiver specifications and product compatibility information.

Technical Tips, Field Notices
Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

OmniSwitch 6900 Series Release 7.3.2.R01 requires 2GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.

OmniSwitch 10K Series Release 7.3.2.R01 requires 4GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.

Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory.

UBoot and FPGA Requirements

The software versions listed below are the minimum required, except where otherwise noted. Switches running the minimum versions, as listed below, do not require any UBoot or FPGA upgrades.

Switches not running the minimum version required should upgrade to the latest UBoot or FPGA that is available with the 7.3.2.R01 AOS software available from Service & Support.

- Newly supported OS6900 models released in 7.3.2 listed in the New Hardware Support section will be factory shipped with the correct Uboot/FPGA. They do not need to be upgraded and should not be downgraded.
- If upgrading from 7.2.1.R02 or higher the Uboot and FPGA should already be at the correct versions listed below. If upgrading from a release prior to 7.2.1.R02 upgrading the Uboot and FPGA according to the table below may be required.
- A separate file containing the Uboot and FPGA upgrade files is available from Service & Support.
- Please refer to the Upgrade Instructions section at the end of these Release Notes for step-by-step instructions on upgrading your switch to support 7.3.2.R01.

OmniSwitch 10K – Release 7.3.2.344.R01 (GA)

Module	Uboot	FPGA
CMM	7.2.1.266.R02[1]	2.0[2]
GNI-C48/U48	7.2.1.266.R02[1]	0.7[2]
GNI-U48 Daughter Card	7.2.1.266.R02[1]	1.4[2]
XNI-U32S	7.2.1.266.R02[1]	2.12[2]
XNI-U16L	7.3.1.325.R01[2]	0.3[2]
XNI-U16E	7.3.1.325.R01[2]	0.3[2]
XNI-U32E	7.3.1.325.R01[2]	0.3[2]
QNI-U4E	7.3.1.325.R01[2]	0.3[2]
QNI-U8E	7.3.1.325.R01[2]	0.3[2]

1. May require upgrade if running an AOS version prior to 7.2.1.R02.
2. Shipped from factory with correct version, no upgrade is available or required.

OmniSwitch 6900-X20/X40 – AOS Release 7.3.2.344.R01(GA)

Hardware	Uboot	FPGA
CMM	7.2.1.266.R02[1]	1.3.0/1.2.0[1]
Expansion Slot	N/A	0.2.0[2]

1. May require upgrade if running an AOS version prior to 7.2.1.R02.
2. Shipped from factory with correct version, no upgrade is available or required.

OmniSwitch 6900-T20/T40 – AOS Release 7.3.2.344.R01(GA)

Hardware	Uboot	FPGA
CMM	7.3.2.134.R01[1]	1.4.0/0.0.0[1]
Expansion Slot	N/A	0.2.0[1]

1. Shipped from factory with correct version, no upgrade is available or required.

AOS Release 7.3.2.R01 Prerequisites

Prior to upgrading to AOS Release 7.3.2.R01 please note the following:

- VRF functionality will be updated to use the new profiles capability in 7.3.2. These new profiles are not compatabile with earlier versions of AOS. It's strongly recommened to create a backup of the 7.3.1 configuration prior to upgrading to prevent the VRF configuration having to be rebuilt if a switch should need to be downgraded.
- A new predefined DCB profile 11 is being introduced in 7.3.2.R01, this will overwrite any existing custom profile 11.

New Hardware Support in 7.3.2.R01

OmniSwitch 6900-T20

10-Gigabit Ethernet (10GBase-T) fixed configuration chassis in a 1U form factor with 20 10-gigabit copper ports, one optional module slot, redundant AC or DC power and front-to-rear or rear-to-front cooling. The switch includes:

- 1 – Console Port (USB Form Factor - RS-232)
- 1 – USB Port (For use with Alcatel-Lucent OS-USB-FLASHDR USB flash drive)
- 1 – EMP Port
- 20 – 10-Gigabit copper ports
- 1 Slot– Optional module
- 1 Slot – Fan Tray
- 2 Slots – Power Supplies (AC or DC)
- Energy Efficient Ethernet
- CAT 5e/6 = 55 meters; CAT 6a/7 = 100 meters
- 1G/10G support

OmniSwitch 6900-T40

10-Gigabit Ethernet (10GBase-T) fixed conguration chassis in a 1U form factor with 40 10-gigabit copper ports, two optional module slots, redundant AC or DC power and front-to-rear or rear-to-front cooling. The switch includes:

- 1 – Console Port (USB Form Factor - RS-232)
- 1 – USB Port (For use with Alcatel-Lucent OS-USB-FLASHDR USB flash drive)
- 1 – EMP Port
- 40 – 10-Gigabit copper ports
- 2 Slots– Optional Modules
- 1 Slot – Fan Tray
- 2 Slots – Power Supplies (AC or DC)
- Energy Efficient Ethernet
- CAT 5e/6 = 55 meters; CAT 6a/7 = 100 meters
- 1G/10G support

OS-XNI-T8

10-Gigabit Ethernet module for the OS6900 series of swicthes with eight 1G/10G copper ports.

New Software Features and Enhancements

The following software features are being introduced with the 7.3.2.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

Features listed as 'Base' are included as part of the base software and do not require any license installation. Features listed as 'Advanced' or "Data Center" require the installation of a license.

7.3.2 New Feature/Enhancements Summary

Feature	Platform	License
Data Center Feature Support		
- FIP Snooping	OS10K/6900	Data Center
- Virtual Maching Performance Monitoring	OS10K/6900	Data Center
Layer 2 Feature Support		
- Dynamic Auto Fabric	OS10K/6900	Base
Layer 3 Feature Support		
- IPv4 over SPBM	OS10K/6900	Advanced
- Interop between PIM & DVMRP	OS10K/6900	Base
- Non-Contiguous Mask and IPv6 Gateway Support	OS10K/6900	Base
- Increase VRF Instances	OS10K/6900	Base
Management/Additional Feature Support		
- Command Abbreviation	OS10K/6900	Base
- Web Services & CLI Scripting	OS10K/6900	Base
- Enhanced Server & Session Limits	OS10K/6900	Base
Additional Feature Support		
- Application Fingerprinting	OS10K/6900	Base
- Fault Propagation and Link Flapping - Wait to Shutdown	OS10K/6900	Base

New Software Features and Enhancements Descriptions

Data Center Feature Descriptions

FIPS Snooping

The OmniSwitch implementation of Fibre Channel over Ethernet (FCoE) Initiation Protocol (FIP) Snooping supports the FCoE technology used to tunnel Fibre Channel frames within Ethernet MAC frames.When the FCoE and FIP snooping functionality is enabled, the OmniSwitch serves as an FCoE transit switch. In this role, the OmniSwitch implementation of Data Center Bridging (DCB) is also used to provide the lossless Ethernet network required to support FCoE.

This implementation of FIP Snooping ensures the security of the FCoE network and maintains a virtual point-to-point network connection between FCoE Nodes (ENodes) and FCoE Forwarder (FCF) devices. An FCoE OmniSwitch is placed between ENodes (servers or other bridges) and a third-party FCF to extend the reach of the FCoE network without extending the physical Fibre Channel (FC) connections.

Virtual Machine Performance Monitoring

The Universal Network Profile (UNP) and Service Assurance Agent (SAA) enhancements described in this section are provided to facilitate the monitoring of Virtual Machine (VM) connectivity across the data center.

UNP/VNP Feature Enhancements

SAA profile to specify jitter and latency thresholds. Assigned to UNP VLAN-based profiles (service-based profiles not supported) to associate these performance monitoring thresholds with the specific UNP.

The Alcatel-Lucent OmniVista network management tool will extract profile information from UNP on the switch and will create SAA sessions based on the UNP profile SAA threshold values. These SAA sessions will operate as regular sessions. When a threshold is reached, a trap is sent to OmniVista, and OmniVista will make the necessary notifications and network modifications. **Note**: Available in a future OV release.

Service Assurance Agent (SAA) Feature Enhancements

- Jitter and round-trip-time (RTT) threshold parameters added to the base SAA.
- SAA analytics; a trap containing the session name is sent when any of the following thresholds are reached:
  - At least one packet lost
  - Warning: Average RTT/Jitter within 10% of threshold
  - Critical: Average RTT/Jitter at or above threshold
- New Shortest Path Bridging (SPB) SAA configuration. SPB will automatically create and start SAA MAC ping sessions for Backbone Edge Bridges (BEBs) based on the SPB SAA configuration.
- SAA results can now be saved to an XML file. Configurable SAA XML parameters enable the switch to periodically store the last five iterations of all SAAs to an XML file on the local switch.

Layer 2 Feature Descriptions

Auto Fabric

The Auto Fabric feature reduces the burden of configuration on the administrator. Dynamic recognition of the neighboring elements will allow for quick, out-of-the box configuration. The focus area for this feature is data center but the feature is applicable in campus LAN environments to reduce administrative overhead. Auto-Fabric can be used to dynamically discover and configure a switch for the LACP, SPB, and MVRP protocols and is supported in both standalone or virtual chassis mode.

Upon boot-up, a system will be a stand alone or part of a Virtual Chassis. If part of a virtual chassis, auto discovery will not operate until after the virtual chassis setup has completed and normal configuration commands are applied from the virtual chassis boot file. The following will then occur:

- The switch will first attempt LACP discovery and auto configuration for a set discovery window.
- After LACP discovery window expires, SPB auto discovery will occur.
- After SPB discovery window expires, MVRP auto discovery will occur.

The LACP and SPB discovred configuration can be saved two ways:

Manually - The configuration discovered by SPB and LACP will be saved to the configuration file (boot.cfg) post discovery if the write memory command is given.

Automatically - The system will save the discovered configuration to the configuration file (boot.cfg) at set periods automatically if this is enabled. For example if LACP discovery is successful, all the LACP configuration will be saved to the configuration file. This feature can be disabled and the interval changed.

Layer 3 Feature Descriptions

IPv4 Over SPBM

The previous implementation of Shortest Path Bridging MAC (SPBM) provides L2 VPN capability that bridges L2 customer LAN segments. Customer edge (CE) devices form peers and exchange routing information, as well as perform the necessary IP forwarding. Then the SPBM Backbone Edge Bridges (BEBs) bridge the already routed IP traffic across the SPBM backbone.

The 7.3.2 release now provides IPv4 over SPBM capability that consolidates the routing functionality of CE devices into BEB devices. The VRF instances on different BEBs are tied together via backbone service instance identifiers (I-SIDs) across the same SPBM backbone that is used to support Layer 2 VPNs.

The OmniSwitch IP over SPBM solution supports two methods for combining Layer 3 routing and SPBM in the same chassis: VPN-Lite and L3 VPN.

- VPN-Lite

  The VPN-Lite method provides a 'gateway' between a regular SPBM service and a router within the same OmniSwitch chassis. This solution provides a specific advantage in that it allows a single box to represent two 'tiers' in a typical 'fat tree' network, which is popular in datacenter deployments.

  In addition, a VPN-Lite configuration can act purely as a L3 VPN when configured correctly. In this mode, existing routing protocols can form adjacencies across the SPBM Provider Backbone Bridge (PBB) network. To keep it purely a L3 VPN, the administrator makes sure that no SPBM Service Access Points (SAPs) that can inject bridged flows are allowed to attach to the VPN's I-SID.

  The VPN-Lite approach uses the SPBM network in the same way a VLAN is used for transporting L3 frames. Each BEB or host can inject frames into the I-SID as needed, and BEBs can decide to bridge our route those frames based on their inner and outer destination MAC address.

- L3 VPN

  When the L3 VPN method is implemented, the OmniSwitch acts as an access or edge router to multiple VRFs and connects these VRFs across an SPBM managed PBB network. This solution is based upon the IETF drafts "IP/IPVPN services with IEEE 802.1aq SPB(B) networks" and uses the proposed IS-IS TLVs to exchange routes between the BEBs that host the same VPN services.

  When the L3 VPN approach is implemented, each VPN is identified by a VRF locally on each BEB and globally in the backbone by an I-SID in the PBB header. SPB IS-IS will import/export routes from the local routing protocols running inside their respective VRFs. In essence, SPB IS-IS is creating tunnels between BEBs through which routed frames are sent to reach their target networks.

  The L3 VPN solution gives an administrator the ability to build VPNs and extend them over a SPBM core without having to define routes and VRFs across that core by hand. The core boxes need only run SPBM.

PIM/DVMRP Interoperability

The OmniSwitch support of interoperability between PIM and DVMRP is based on rules defined in RFC 2715 and Multicast Border Router (MBR) functionality defined in the PIM-SM specification (RFC 4601). The supported MBR functionality allows receivers and sources within PIM and DVMRP domains to communicate and satisfy RFC 2715 rules.

MBR functionality is configured and enabled on OmniSwitches that are located at points where PIM and DVMRP regions interconnect. An MBR first pulls down packets generated within the PIM domain and injects them into the DVMRP domain. Then the MBR imports packets generated within the DVMRP domain so that

they can be delivered to group members inside the PIM domain, using PIM mechanisms. In the case of transit networks, the MBR acts to pass the multicast traffic through both the PIM and DVMRP domains.

The MBR functionality implemented for the OmniSwitch supports interoperability between a PIM and DVMRP domain. Interoperability between PIM and other protocols or between multiple PIM domains is not supported. In addition, PIM support refers only to PIM-DM and PIM-SM (PIM-SSM is not supported).

Non-contiguous Mask and IPv6 Gateway Support

This section describes a new feature for the OmniSwitch 6900 and 10K platforms that expands the accepted inputs for the Access Control List (ACL) netmask to facilitate load distribution through Policy Based Routing (PBR). The feature allows masks consisting of any combination of 0s and 1s. Prior to this change only traditional netmasks were supported and only allowed up to 8 bits of 0 to be sparsely distributed in the mask. This new feature supports both IPv4 and IPv6 non-contiguous address masks in policy condition statements that contain any sequence of 0 and 1 bits. Additionally, permanent gateway support has been enhanced to provide the ability to forward to an IPv6 gateway address.

Increase VRF Instances

To increase the number of supported VRFs, 7.3.2 introduces VRF profiles. Two types of profiles can now be configured called ‘low’ and ‘max’. A ‘max’ profile is the same as a VRF configured prior to 7.3.2 with no restrictions on the number of IPv4 protocols.

The ‘low’ profile VRF restricts all routing protocols and provides support only for static routes and routes imported from other VRFs.

- OS10K – Supports 512 low profiles or 64 max profiles
- OS6900 – Supports 128 low profiles or 64 max profiles

Note: When mixing low and max profiles the total number of each type will be dependant upon the available sytem resources.

Management Feature Descriptions

Command Abbreviation

Allow users to enter abbreviated commands in the CLI for the command to be accepted. This input only works when enough characters of a keyword are entered to completely identify a single branch of the options available under the preceding keyword. Only pure CLI keywords are auto-completed; Bash or Linux keywords or command names such as “ls” or “awk” are not completed.

Ex. “show vlan” can be abbreviated to “sh vl”

Web Services

The Web Services feature provides the ability to customize and extend the management interface on AOS devices. It supports the use of CLI scripting in AOS as well as a REST based 'web' interface that interacts with AOS management variables (MIB) and CLI commands. It provides two methods for configuration through either the direct handling of MIB variables or the use of CLI commands and supports both XML and JSON response formats.

An example Python library has been created which can be used by any Python Consumer communicating with the AOS Web Services. The library is available in source form and provides a tool allowing developers to learn how to write code that communicates with the OmniSwitch Web Services. In addition, this library can also be used as a standalone query tool using the command line.

Additional Feature Descriptions

Application Fingerprinting

The OmniSwitch Application Fingerprinting feature attempts to detect and identify remote applications by scanning the payload of IP packets and comparing the payload to pre-defined bit patterns (application signatures). Once an application is identified, Application Fingerprinting collects and stores information about the application flow in a database on the local switch. Additional configurable options for this feature include the ability to apply QoS policy list rules to the identified flow and generating SNMP traps when a signature match occurs.

Application Fingerprinting operates in three modes:

- Monitoring – No action, accounting and visibility
- QoS – If there is a match, applies QoS policy on port
- UNP – If match, applies QoS on only if source MAC/VLAN matches UNP profile

Fault Propagation and Link Monitoring

The Link Monitoring feature is used to monitor interface status to minimize the network protocol reconvergence that can occur when an interface becomes unstable. To track the stability of an interface, this feature monitors link errors and link flaps during a configured timeframe. If the number of errors or link flaps exceeds configured thresholds during this time frame, the interface is shut down.

There are no explicit Link Monitoring commands to recover a port from a Link Monitoring shutdown. Such ports are subject to the interfaces violation recovery mechanisms configured for the switch. The following capabilites are provided:

- Wait to Restore Time – Introduces a delay before the interface becomes operational allowing the network to convergence more gracefully.
- Interface errors monitoring - Physical errors such as CRC, Lost frames, Errors frames and Alignment errors are monitored. When excessive errors are detected, the interface will be shutdown.
- Interface flapping - When excessive interface flapping is detected, the interface will be shutdown.
- Permanent shutdown - When an interface has been shutdown too many times it can be placed in this mode requiring it to be enabled administratively.

Link Fault Propagation

The Link Fault Propagation (LFP) feature provides a mechanism to propagate a local interface failure into another local interface. In many scenarios, a set of ports provide connectivity to the network. If all these ports go down, the connectivity to the network is lost. However, the remote end remains unaware of this loss of connectivity and continues to send traffic that is unable to reach the network. To solve this problem, LFP does the following:

- Monitors a group of interfaces (configured as source ports).

o If all the source ports in the group go down, LFP waits a configured amount of time then shuts down another set of interfaces (configured as destination ports) that are associated with the same group.

o When any one of the source ports comes back up, all of the destination ports are brought back up and network connectivity is restored.

Interface Violation Recovery
The OmniSwitch allows features to shutdown an interface when a violation occurs on that interface. To support this functionality, the following interfaces violation recovery mechanisms are provided:

o Manual recovery of a downed interface using a CLI command.

o An automatic recovery timer that indicates how much time a port remains shut down before the switch automatically brings the port back up.

o A maximum number of recovery attempts setting that specifies how many recoveries can occur before a port is permanently shutdown

o A wait-to-restore timer that indicates the amount of time the switch waits to notify features that the port is back up.

o An SNMP trap that is generated each time an interface is shutdown by a feature. This can occur even when the interface is already shutdown by another feature. The trap also indicates the reason for the violation

o An SNMP trap that is generated when a port is recovered. The trap also includes information about how the port was recovered.

Ethernet OAM
This feature is used to propagate OAM Connectivity Fault Management (CFM) events into the interface that is attached to a MEP.

Wait to Shutdown
The wait-to-shutdown (WTS) timer is used to implement a delay before an interface is made non-operational for other applications such as data, control and management. Only after the timer has expired will the interface become disabled allowing network protocols to converge more gracefully. The timer value is configured on a per-port basis and is started whenever one of the following link-up events occurs:

An interface is administratively brought down.

An interface is shutdown from a violation.

An interface is made operationally down when the cable is unplugged in.

Enhanced Server and Session Limits

This following server and session limits have been increased as described below:

- The number of concurrent Telnet sessions has been increased to 6.
- The maximum number of Syslog servers supported has been increased from 4 to 12.
- The maximum number of NTP servers supported has been increased from 2 to 12.
- The maximum number of characters for system name is now 32.
- The maximum number OSPFv3 neighbors is now 128.

  Note: Syslog servers can be added to any VRF up to a maximum of 12. Each Syslog server will receive the same information, the server can be used to filter the information if required.

SNMP Traps

The following table provides a list of SNMP traps managed by the switch.

No.	Trap Name	Platforms	Description
0	coldStart	OS10K OS6900	The SNMP agent in the switch is reinitiating and itsk configuration may have been altered.
1	warmStart	OS10K OS6900	The SNMP agent in the switch is reinitiating itself and its configuration is unaltered.
2	linkDown	OS10K OS6900	The SNMP agent in the switch recognizes a failure in one of the communications links configured for the switch.
3	linkUp	OS10K OS6900	The SNMP agent in the switch recognizes that one of the communications links configured for the switch has come up.
4	authenticationFailure	OS10K OS6900	The SNMP agent in the switch has received a protocol message that is not properly authenticated.
5	entConfigChange	OS10K OS6900	An entConfigChange notification is generated when a conceptual row is created, modified, or deleted in one of the entity tables.
6	policyEventNotification	OS10K OS6900	The switch notifies the NMS when a significant event happens that involves the policy manager.
7	chassisTrapsStr	OS10K OS6900	A software trouble report (STR) was sent by an application encountering a problem during its execution.
8	chassisTrapsAlert	OS10K OS6900	A notification that some change has occurred in the chassis.
9	chassisTrapsStateChange	OS10K OS6900	An NI status change was detected.
10	chassisTrapsMacOverlap	OS10K OS6900	A MAC range overlap was found in the backplane eeprom.
11	vrrpTrapNewMaster	OS10K OS6900	The SNMP agent has transferred from the backup state to the master state.
12	vrrpTrapAuthFailure	OS10K OS6900	This trap is not supported.
13	healthMonModuleTrap	OS10K OS6900	Indicates a module-level threshold was crossed.

No.	Trap Name	Platforms	Description
14	healthMonPortTrap	OS10K OS6900	Indicates a port-level threshold was crossed.
15	healthMonCmmTrap	OS10K OS6900	This trap is sent when the Module-level rising/falling threshold is crossed.
16	bgpEstablished	OS10K OS6900	The BGP routing protocol has entered the established state.
17	bgpBackwardTransition	OS10K OS6900	This trap is generated when the BGP router port has moved from a more active to a less active state.
18	esmDrvTrapDropsLink	OS10K OS6900	This trap is sent when the Ethernet code drops the link because of excessive errors.
19	portViolationTrap	OS10K OS6900	This trap is sent when a port violation occurs. The port violation trap will indicate the source of the violation and the reason for the violation.
20	dvmrpNeighborLoss	OS10K OS6900	A 2-way adjacency relationship with a neighbor has been lost. This trap is generated when the neighbor state changes from “active” to “one-way,” “ignoring” or “down.” The trap is sent only when the switch has no other neighbors on the same interface with a lower IP address than itself.
21	dvmrpNeighborNotPruning	OS10K OS6900	A non-pruning neighbor has been detected in an implementation-dependent manner. This trap is generated at most once per generation ID of the neighbor. For example, it should be generated at the time a neighbor is first heard from if the prune bit is not set. It should also be generated if the local system has the ability to tell that a neighbor which sets the prune bit is not pruning any branches over an extended period of time. The trap should be generated if the router has no other neighbors on the same interface with a lower IP address than itself.
22	risingAlarm	OS10K OS6900	An Ethernet statistical variable has exceeded its rising threshold. The variable’s rising threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
23	fallingAlarm	OS10K	An Ethernet statistical variable has dipped below its falling threshold. The variable’s

No.	Trap Name	Platforms	Description
		OS6900	falling threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
24	stpNewRoot	OS10K OS6900	Sent by a bridge that became the new root of the spanning tree.
25	stpRootPortChange	OS10K OS6900	A root port has changed for a spanning tree bridge. The root port is the port that offers the lowest cost path from this bridge to the root bridge.
26	mirrorConfigError	OS10K OS6900	This trap is sent when any NI fails to configure mirroring. Due to this error, port mirroring session will be terminated.
27	mirrorUnlikeNi	OS10K OS6900	The mirroring configuration is deleted due to the swapping of different NI board type. The Port Mirroring session which was active on a slot cannot continue with the insertion of different NI type in the same slot.
28	slbTrapOperStatus	OS10K OS6900	A change occurred in the operational status of the server load balancing entity.
29	sessionAuthenticationTrap	OS10K OS6900	An authentication failure trap is sent each time a user authentication is refused.
30	trapAbsorptionTrap	OS10K OS6900	The absorption trap is sent when a trap has been absorbed at least once.
31	alaDoSTrap	OS10K OS6900	Indicates that the sending agent has received a Denial of Service (DoS) attack.
32	ospfNbrStateChange	OS10K OS6900	Indicates a state change of the neighbor relationship.
33	ospfVirtNbrStateChange	OS10K OS6900	Indicates a state change of the virtual neighbor relationship.
34	lnkaggAggUp	OS10K OS6900	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.
35	lnkaggAggDown	OS10K OS6900	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
36	lnkaggPortJoin	OS10K	This trap is sent when any given port of the link aggregate group goes to the

No.	Trap Name	Platforms	Description
		OS6900	attached state.
37	lnkaggPortLeave	OS10K OS6900	This trap is sent when any given port detaches from the link aggregate group.
38	lnkaggPortRemove	OS10K OS6900	This trap is sent when any given port of the link aggregate group is removed due to an invalid configura tion.
39	monitorFileWritten	OS10K OS6900	This trap is sent when the amount of data requested has been written by the port monitoring instance.
40	alaVrrp3TrapProtoError	OS10K OS6900	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
41	alaVrrp3TrapNewMaster	OS10K OS6900	The SNMP agent has transferred from the backup state to the master state.
42	chassisTrapsPossibleDuplicateMac	OS6900	The old PRIMARY element cannot be detected in the stack. There is a possiblity of a duplicate MAC address in the network
43	lldpRemTablesChange	OS10K OS6900	A lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes.
44	pimNeighborLoss	OS10K OS6900	A pimNeighborLoss notification signifies the loss of an adjacency with a neighbor.
45	pimInvalidRegister	OS10K OS6900	An pimInvalidRegister notification signifies that an invalid PIM Register message was received by this device
46	pimInvalidJoinPrune	OS10K OS6900	A pimInvalidJoinPrune notification signifies that an invalid PIM Join/Prune message was received by this device.
47	pimRPMappingChange	OS10K OS6900	An pimRPMappingChange notification signifies a change to the active RP mapping on this device.
48	pimInterfaceElection	OS10K OS6900	An pimInterfaceElection notification signifies that a new DR or DR has been elected on a network.
49	pimBsrElectedBSRLostElection	OS10K OS6900	This trap is sent when the current E-BSR loses an election to a new Candidate-BSR.
50	pimBsrCandidateBSRWinElection	OS10K OS6900	This trap is sent when a C-BSR wins a BSR Election.
51	lpsViolationTrap	OS10K	A Learned Port Security (LPS) violation has

No.	Trap Name	Platforms	Description
		OS6900	occurred.
52	lpsPortUpAfterLearningWindowExpiredT	OS10K OS6900	When an LPS port joins or is enabled after the Learning Window is expired, the MAC address learning on the port will be disabled, and this trap is generated as a notification.
53	lpsLearnMac	OS10K OS6900	Generated when an LPS port learns a bridged MAC.
54	gvrpVlanLimitReachedEvent	OS10K OS6900	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
55	alaNetSecPortTrapAnomaly	OS10K OS6900	Trap for an anomaly detected on a port.
56	alaNetSecPortTrapQuarantine	OS10K OS6900	Trap for an anomalous port quarantine.
57	ifMauJabberTrap	OS10K OS6900	This trap is sent whenever a managed interface MAU enters the jabber state.
58	udldStateChange	OS10K OS6900	Generated when the state of the UDLD protocol changes.
59	ndpMaxLimitReached	OS10K OS6900	This IPv6 Trap is sent when the hardware table has reached the maximum number of entries supported.
60	ripRouteMaxLimitReached	OS10K OS6900	This trap is sent when the RIP database reaches the supported maximum number of entries. When the maximum number is reached, RIP discards any new updates.
61	ripngRouteMaxLimitReached	OS10K OS6900	This trap is sent when the RIPng database reaches the supported maximum number of entries. When the maximum number is reached, RIPng discards any new updates.
62	alaErpRingStateChanged	OS10K OS6900	This trap is sent when the ERP Ring State has changed from “Idle” to “Protection”.
63	alaErpRingMultipleRpl	OS10K OS6900	This trap is sent when multiple RPLs are detected in the Ring.
64	alaErpRingRemoved	OS10K OS6900	This trap is sent when the Ring is removed dynamically.
65	ntpMaxAssociation	OS10K OS6900	This trap is generated when the maximum number of peer and client associations configured for the switch is exceeded.

No.	Trap Name	Platforms	Description
66	ddmTemperatureThresholdViolated	OS10K OS6900	This trap is sent when an SFP/ XFP/SFP+ temperature has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/ XFP/SFP+ temperature.
67	ddmVoltageThresholdViolated	OS10K OS6900	This trap is sent when SFP/XFP/ SFP+ supply voltage has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ supply voltage.
68	ddmCurrentThresholdViolated	OS10K OS6900	This trap is sent when if an SFP/ XFP/SFP+ Tx bias current has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx bias current.
69	ddmTxPowerThresholdViolated	OS10K OS6900	This trap is sent when an SFP/ XFP/SFP+ Tx output power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx output power.
70	ddmRxPowerThresholdViolated	OS10K OS6900	This trap is sent when an SFP/ XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
71	webMgtServerErrorTrap	OS10K OS6900	This trap is sent to management station(s) when the Web Management server goes into error state after becoming unreachable twice within a minute.
72	multiChassisIpcVlanUp	OS10K OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is up.
73	multiChassisIpcVlanDown	OS10K OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is down.
74	multiChassisMisconfigurationFailure	OS10K OS6900	This trap is sent to indicate a mis-configuration due to Chassis Id or IPC VLAN.
75	multiChassisHelloIntervalConsisFailure	OS10K OS6900	This trap is sent to indicate a hello interval consistency failure.

No.	Trap Name	Platforms	Description
76	multiChassisStpModeConsisFailure	OS10K OS6900	This trap is sent to indicate an STP mode consistency failure.
77	multiChassisStpPathCostModeConsisFailure	OS10K OS6900	This trap is sent to indicate an STP path cost mode consistency failure.
78	multiChassisVflinkStatusConsisFailure	OS10K OS6900	This trap is sent to indicate a VFLink status consistency failure.
79	multiChassisStpBlockingStatus	OS10K OS6900	This trap is sent to indicate the STP status for some VLANs on the VFLink is in blocking state.
80	multiChassisLoopDetected	OS10K OS6900	This trap is sent to indicate a loop has been detected.
81	multiChassisHelloTimeout	OS10K OS6900	This trap is sent to indicate the hello timeout has occurred.
82	multiChassisVflinkDown	OS10K OS6900	This trap is sent to indicate the VFLink is down.
83	multiChassisVFLMemberJoinFailure	OS10K OS6900	This trap is sent to indicate a port configured as virtual-fabric member is unable to join the virtual-fabric link.
84	alaDHLVlanMoveTrap	OS10K OS6900	When linkA or linkB goes down or comes up and both ports are are part of some vlan-map, this trap is sent to the Management Entity, with the DHL port information.
85	alaDhcpClientAddressAddTrap	OS10K OS6900	This trap is sent when a new IP address is assigned to DHCP Cli-ent interface.
86	alaDhcpClientAddressExpiryTrap	OS10K OS6900	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address.
87	alaDhcpClientAddressModifyTrap	OS10K OS6900	This trap is sent when the DHCP client is unable to obtain the existing IP address and a new IP address is assigned to the DHCP client
88	vRtrIsisDatabaseOverload	Not supported	This notification is generated when the system enters or leaves the overload state.
89	vRtrIsisManualAddressDrops	Not supported	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
90	vRtrIsisCorruptedLSPDetected	Not supported	This notification is generated when an LSP that was stored in memory has become corrupted.

No.	Trap Name	Platforms	Description
91	vRtrIsisMaxSeqExceedAttempt	Not supported	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
92	vRtrIsisIDLenMismatch	Not supported	A notification sent when a PDU is received with a different value of the System ID Length.
93	vRtrIsisMaxAreaAddrsMismatch	Not supported	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
94	vRtrIsisOwnLSPPurge	Not supported	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
95	vRtrIsisSequenceNumberSkip	Not supported	When an LSP is received without a System ID and different contents.
96	vRtrIsisAutTypeFail	Not supported	A notification sent when a PDU is received with the wrong authentication type field.
97	vRtrIsisAuthFail	Not supported	A notification sent when a PDU is received with an incorrent authentication information field.
98	vRtrIsisVersionSkew	Not supported	A notification sent when a a Hello PDU is received from an IS running a different version of the protocol.
99	vRtrIsisAreaMismatch	Not supported	A notification sent when a Hello PDU is received from an IS which does not share any area address.
100	vRtrIsisRejectedAdjacency	Not supported	A notification sent when a Hello PDU is received from an IS, but does not establish an adjacency due to a lack of resources.
101	vRtrIsisLSPTooLargeToPropagate	Not supported	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
102	vRtrIsisOrigLSPBufSizeMismatch	Not supported	A notification sent when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for the originating L1LSP BufferSize or originating L2LSPBufferSize respectively. Also when a Level 1 LSP or Level2 LSP is received containing the originating LSPBufferSize option and the value in the PDU option field does not match the local value for originating L1LSP BufferSize or originatingL2LSP BufferSize respectively.
103	vRtrIsisProtoSuppMismatch	Not supported	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.

No.	Trap Name	Platforms	Description
104	vRtrIsisAdjacencyChange	Not supported	A notification sent when an adjacency changes state, entering or leaving state up. The first 6 bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
105	vRtrIsisCircIdExhausted	Not supported	A notification sent when ISIS cannot be started on a LAN interface because a unique circId could not be assigned due to the exhaustion of the circId space.
106	vRtrIsisAdjRestartStatusChange	Not supported	A notification sent when an adjancency's graceful restart status changes.
107	mvrpVlanLimitReachedEvent	OS10K OS6900	This trap is sent when the num-ber of VLANs learned dynamically by MVRP reaches the configured limit.
108	alaHAVlanClusterPeerMismatch	OS10K OS6900	The trap is sent when parameteras configured for this cluster ID (Level 1 check) does not match accross the MCLAG peers.
109	alaHAVlanMCPeerMismatch	OS10K OS6900	The trap is sent when when the cluster parameters are matching on the peers, but MCLAG is not configured or clusters are not in operational state.
110	alaHAVlanDynamicMAC	OS10K OS6900	The trap is sent when the dynamic MAC is learned on non-server cluster port
111	unpMcLagMacIgnored	OS10K OS6900	This trap is sent when a MAC/User is dropped because the VLAN does not exist or UNP is not enabled on the MCLAG.
112	unpMcLagConfigInconsistency	OS10K OS6900	This trap is sent when a configuration becomes “Out of Sync".
113	multiChassisGroupConsisFailure	OS10K OS6900	This trap is sent when there is an inconsistency between local and peer chassis group.
114	multiChassisTypeConsisFailure	OS10K OS6900	This trap is sent when there is an inconsistency between local and peer chassis group.
115	alaPimNonBidirHello	OS10K OS6900	This trap is sent when a bidir-capable router has received a PIM hello from a non-bidir-capable router. It is generated whenever the counter alaPimsmNon-BidirHelloMsgsRcvd is incremented, subject to the rate limit specified by alaPimsmNonBidirHelloNotificationPeriod.
116	dot1agCfmFaultAlarm	OS10K	This trap is sent when a MEP has a persistent defect condition. A notification

No.	Trap Name	Platforms	Description
		OS6900	(fault alarm) is sent to the management entity with the OID of the MEP that has detected the fault.
117	alaSaaIPIterationCompleteTrap	OS10K OS6900	This trap is sent when an IP SAA iteration is completed.
118	alaSaaEthIterationCompleteTrap	OS10K OS6900	This trap is sent when when an eth-LB or Eth-DMM SAA iteration is completed.
119	alaSaaMacIterationCompleteTrap	OS10K OS6900	This trap is sent when a MAC iteration is complete.
120	virtualChassisStatusChange	OS10K OS6900	This trap is sent when a chassis status change is detected.
121	virtualChassisRoleChange	OS10K OS6900	This trap is sent when a chassis role change is detected.
122	virtualChassisVflStatusChange	OS10K OS6900	This trap is sent when s vflink status change is detected.
123	virtualChassisVflMemberPortStatusCh	OS10K OS6900	This trap is sent when a vflink member port has a change of status.
124	virtualChassisVflMemberPortJoinFail	OS10K OS6900	This trap is sent when a port configured as virtual-fabric member is unable to join the virtual-fabric link.
125	lldpRemTablesChange	OS10K OS6900	This trap is sent when the value of lldpStatsRemTablelastChange Time changes. It can be utilized by an NMS to trigger LLDP remote systems table maintenance polls.
126	vRtrLdpInstanceStateChange	OS10K OS6900	This trap is sent when the LDP module changes state either administratively or operationally.
127	evbFailedCdcpTlvTrap	OS10K OS6900	This trap is sent when bridge receives a CDCP packet with: - Wrong TLV type, or - Wrong OUI, or - Role is set to Bridge, or - Wrong default channel(scid), or - Incorrect channel number(scid).
128	evbFailedEvbTlvTrap	OS10K	This trap is sent when bridge receives an

No.	Trap Name	Platforms	Description
		OS6900	EVBTLV packet with: - Wrong TLV type. or - Incorrect TLV length, or - Wrong OUI.
129	evbUnknownVsiManagerTrap	OS10K OS6900	This trap is sent when bridge receives a VDP packet with: - Unknown Manager ID type, or - Wrong Manager ID length.
130	evbVdpAssocTlvTrap	OS10K OS6900	This trap is sent when bridge receives an ASSOC TLV in a VDP packet with: - Null VID found and number of entry field is not 1, or - Unknown filter format, - Null VID on De-Assoc TLV type, or - VSI included more than Max number of filter info entries
131	evbCdcpLldpExpiredTrap	OS10K OS6900	This trap is sent when an LLDP Timer expired in bridge. The timer expires when LLDP doesn't not receive CDCP TLV within a specified interval.
132	evbTlvExpiredTrap	OS10K OS6900	This trap is sent when an LLDP Timer expired in bridge. The timer expires when LLDP doesn't not receive EVB TLV within a specified interval.
133	evbVdpKeepaliveExpiredTrap	OS10K OS6900	This trap is sent when a VDP Keep Alive Timer expired in bridge. The timer expires when the bridge doesn't not receive VDP Keep Alive message within a specified interval.
134	smgrServiceError	OS10K OS6900	This trap is sent when there is a failure to create/delete a service.
135	smgrServiceHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a service, or to program the hardware tables for a service.
136	smgrServiceSapError	OS10K OS6900	This trap is sent when there is a failure to create/delete a Service Access Point.
137	smgrServiceSapHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a SAP, or to program the hardware tables for a SAP.

No.	Trap Name	Platforms	Description
138	smgrServiceSdpError	OS10K OS6900	This trap is sent when there is a failure to create/delete a Service Distribution Point.
139	smgrServiceSdpHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP, or to program the hardware tables for an SDP.
140	smgrServiceSdpBindError	OS10K OS6900	This trap is sent when there is a failure to create/delete an SDP Bind.
141	smgrServiceSdpBindHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP Bind, or to program the hardware tables for an SDP Bind.
142	smgrGeneralError	OS10K OS6900	This trap is sent when there is a .general system failure detected during normal system operation.
143	smgrStatusChange	OS10K OS6900	This trap is sent when there is a status change for a group of selected services.
144	portViolationNotificationTrap	OS10K OS6900	This trap is sent when a port violation is cleared.
145	multiChassisConsisFailureRecovered	OS10K OS6900	This trap is sent when the system has recovered from a multi-chassis inconsistency between the local and peer switches
146	alaSaaPacketLossTrap	OS10K OS6900	This trap is sent when a a packet is lost during a test.
147	alaSaaJitterThresholdYellowTrap	OS10K OS6900	This trap is sent when the Jitter Threshold crosses 90%.
148	alaSaaRTTThresholdYellowTrap	OS10K OS6900	This trap is sent when the RTT Threshold crosses 90%.
149	alaSaaJitterThresholdRedTrap	OS10K OS6900	This trap is sent when the Jitter threshold is crossed.
150	alaSaaRTTThresholdRedTrap	OS10K OS6900	This trap is sent when the RTT threshold is crossed.
151	chassisTrapsDuplicateMacClear	OS10K OS6900	This trap is sent when the old Master Chassis has rejoined the Virtual Chassis as a slave.
152	alaFipsConfigFilterResourceLimit	OS10K OS6900	The allowed maximum percentage of filter resources configured from the allocated FIPS resources is exceeded.

No.	Trap Name	Platforms	Description
153	virtualChassisUpgradeComplete	OS10K OS6900	Critical trap indicates whether the software upgrade process has failed after a timeout or completed successfully. Note that if the process fails, it may be still possible for the system to recover if the process successfully completes later after the expired timeout.
154	appFPSignatureMatchTrap	OS10K OS6900	This trap is sent when a traffic flow matches an application signature.

Unsupported Software Features

The following CLI commands and Web Management options may be available in the switch software for the following features. These features are not supported:

Feature	Platform	License
Dual-Home Link Aggregation	OS10K/OS6900	Base
NetSec	OS10K/OS6900	Base

Unsupported CLI Commands

The following CLI commands may be available in the switch software for the following features. These commands are not supported:

Software Feature	Unsupported CLI Commands
Qos	show qos qsi wred-stats (OS10K)
Source Learning	mac-learning mode [distributed | centralized]
Chassis	reload slot
SLB	server-cluster port all

Open Problem Reports and Feature Exceptions

The problems listed here include problems known at the time of the product's release. Any problems not discussed in this section should be brought to the attention of the Alcatel-Lucent Technical Support organization as soon as possible. Please contact customer support for updates on problem reports (PRs) where no known workaround was available at the time of release.

FIP Snooping

PR	Description	Workaround
180621	If FIP snooping is disabled and then enabled on an edge switch the FCoE sessions may be disrupted.	Edge switches directly connected to Enodes must relogin for FCoE sessions to re-enable.

Hardware

PR	Description	Workaround
179832	Upon bootup of an OS6900-T20/T40 or XNI-T8 module a burst of CRC errors is sometimes recorded in the interface stats counter.	There is no known workaround
179315	On an OS6900-T20/T40 or XNI-T8 module some CRC errors are seen when ports are running at 1G speed.	When negotiating an OS6900 10Gbase-T port down to 1000Base-T speed and interfacing with legacy 1000Base-T devices that do not comply to the maximum output droop requirements in 802.3 clause 40.6.1.2.2 can result in an error rate exceeding 10e-12. The link partner must be able to accommodate a minimum Open Circuit Inductance of 100uH.

Multicast

PR	Description	Workaround
179687	When DVMRP default route on an MBR is disabled, flow stops immediatly, but source list continues to remain upto 1 min (against 30 secs source timeout). IPMS source-timeout value, which is 30 seconds by default, and the fact that BCD currently uses a 30 second timeout of it's own for aging the multicast flows. So you should expect the flow to actually get aged out any where from the ipms source-timeout value plus up to 30 additional seconds for the BCD timeout. This is what is causing the flow to take up to 1 minute to actually get aged out.	There is no known workaround at this time. IPMS source-timeout value, which is 30 seconds by default, and the BCD currently uses a 30 second timeout for aging the multicast flows. The flow will age out between the IPMS source-timeout value plus up to 30 additional seconds for the BCD timeout.

SPB

PR	Description	Workaround
180244	The number of SPB/L3VPN routes that can be injected into ISIS by a BEB node is constrained by the size/available space in the BEB's self-generated LSP. In the most optimistic circumstance (minimal local service configuration, etc), at most approximately 30K L3VPN routes can be advertised by a given BEB. As the ISIS configuration grows in size (which increases consumption of LSP space), fewer and fewer L3VPN prefixes fit into the remaining LSP space. When LSP space is insufficient to carry all locally-generated LSP data (which includes L3VPN routes), ISIS enters the overload state.	Reduce LSP space requirements by either decreasing the number of injected L3VPN routes or condensing/reducing other areas of the configuration (services, bvlans, neighbors, etc).

Virtual Chassis

PR	Description	Workaround
177992	After entering the 'reload all' command on an OS10K in a Virtual Chassis, cmmB will sometimes become the primary CMM.	There is no known workaround at this time.
180321	In a virtual chassis configuration when a large VRF configuration is applied through the "configuration apply …" command, some L3 traffic may not be forwarded properly after a virtual chassis management failover.	Use CLI command to create/modify VRF config on run time instead of snapshot file using "configuration apply" CLI command.
180572	Some traffic loss may be seen on VFL link when sending at wire rate. Since all packets that traverse the VFL have an additional 16 byte header prepended to the packet this reduces the effective bandwidth of a given VFL port. There is a slight compensation due to interpacket gap being smaller when sending these packets.	There is no workaround for this issue

Webview

PR	Description	Workaround
178308	After a virtual chassis takeover Webview does not display the DDM information for transceivers on the new Master chassis.	Use the 'show interfaces ddm' CLI command

Hot Swap/Redundancy Feature Guidelines

Hot Swap Feature Guidelines

- Hot swap of like modules is supported.
- Hot swap of unlike modules is not supported.
- Hot insertion, the insertion of a module into a previously empty slot, is supported on 10-Gigabit modules (i.e. OS-XNI-U4 and OS-XNI-U12).
- Hot insertion, the insertion of a module into a previously empty slot, is not supported on 40-Gigabit modules (i.e. OS-QNI-U3 and OS-HNI-U6) due to the hardware having to be reset for 40-Gigabit support. After hot-inserting a 40-Gigabit module, a reboot is required.
- For the OS6900-X40 wait for first module to become operational before adding the second module.

Hot Swap Procedure

The following steps must be followed when hot-swapping expansion modules.

1. Disconnect all cables from transceivers on module to be hot-swapped.
2. Extract all transceivers from module to be hot-swapped.
3. Extract the module from the chassis and wait approximately 30 seconds before inserting replacement.
4. Insert replacement module of same type.
5. Wait for a message similar to the following to display on the console or issue the -> show module status command and wait for operational status to show 'UP':

   ChassisSupervisor niMgr info message:

   +++ Expansion module 2 ready!
6. Re-insert all transceivers into new module.
7. Re-connect all cables to transceivers.

Hot Swap Time Guidelines

- All module extractions must have a 30 second interval before initiating another hot swap activity.
- All module insertions must have a 5 minute interval AND the OK2 LED blinking green before initiating another hot swap activity.

Technical Support

Alcatel-Lucent technical support is committed to resolving our customer's technical issues in a timely manner. Customers with inquiries should contact us at:

Region	Phone Number
North America	800-995-2696
Latin America	877-919-9526
Europe Union	+800 00200100 (Toll Free) or +1(650)385-2193
Asia Pacific	+65 6240 8484

Email: esd.support@alcatel-lucent.com

Internet: Customers with Alcatel-Lucent service agreements may open cases 24 hours a day via Alcatel-Lucent's support web page at: service.esd.alcatel-lucent.com.

Upon opening a case, customers will receive a case number and may review, update, or escalate support cases on-line. Please specify the severity level of the issue per the definitions below. For fastest resolution, please have telnet or dial-in access, hardware configuration—module type and revision by slot, software revision, and configuration file available for each switch.

Severity 1 Production network is down resulting in critical impact on business—no workaround available.

Severity 2 Segment or Ring is down or intermittent loss of connectivity across network.

Severity 3 Network performance is slow or impaired—no loss of connectivity or data.

Severity 4 Information or assistance on product feature, functionality, configuration, or installation.

Appendix A: Upgrading an OmniSwitch to 7.3.2.R01

Overview

These instructions document how to upgrade the following OmniSwitch products to 7.3.2.R01 AOS software. Release 7.3.2.R01 is supported on the OS10K and OS6900 switches.

Prerequisites

This instruction sheet requires that the following conditions exist, or are performed, before upgrading. The person performing the upgrade must:

- be the responsible party for maintaining the switch's configuration
- be aware of any issues that may arise from a network outage caused by improperly loading this code
- understand that the switch must be rebooted and network users will be affected by this procedure
- have a working knowledge of the switch to configure it to accept an FTP connection through the EMP or Network Interface (NI) Ethernet port
- Read the 7.3.2.R01 GA Release Notes prior to performing any upgrade for information specific to this release.
- Ensure there is a current certified configuration on the switch so that the upgrade can be rolled-back if required.
- Verify the current versions of Uboot on the OS10K. If they meet the minimum requirements, (i.e. they were already upgraded during the 7.2.1.R02 upgrade) then only an upgrade of the AOS images is required.
- Verify the current versions of Uboot and FPGA on the OS6900. If they meet the minimum requirements, (i.e. they were already upgraded during the 7.2.1.R02 upgrade) then only an upgrade of the AOS images is required.

Do not proceed until all the above prerequisites have been met. Any deviation from these upgrade procedures will result in the malfunctioning of the switch. All steps in these procedures should be reviewed before beginning.

Note: The examples below use the '**working'** directory as the upgrade directory, however any user-defined directory can be used for the upgrade.

OmniSwitch 10K – Upgrade Instructions

Upgrading OS10K Switches to 7.3.2.R01 consists of the following steps. The steps should be performed in order:

1. Download the upgrade files

Go the to Alcatel-Lucent Service and Support Website and download and unzip the upgrade files. The Zip File contains the following files:

- Image Files - Reni.img, Ros.img (7.3.2.R01)
- U-Boot File - u-boot.7.2.1.R02.266.tar.gz (if required)

2. FTP the upgrade files to the switch

FTP the upgrade files to the following directories on the **Primary CMM** of the switch you are upgrading:

- Image Files - Reni.img, Ros.img - **/flash/working** directory
- U-Boot File - u-boot.7.2.1.R02.266.tar.gz - /**flash** directory

3. Upgrading the U-Boot File (if required)

Follow the steps below to upgrade the U-Boot File on the CMM(s) and NI(s).

CMM Upgrade

Follow the steps below to upgrade the U-Boot File on the CMM(s). If you have dual CMMs, you must update the U-Boot File on both CMMs.

1. Execute the **update uboot cmm slot** command to update the U-Boot File on the Primary CMM (CMM A). The command below is used if the Primary CMM is in Slot 1 ("cmm 1"). If the Primary CMM is in Slot 2, enter "cmm 2".

```
OS10K-> update uboot cmm 1 file u-boot.7.2.1.R02.266.tar.gz
```

2. If you are updating a single CMM, the process is complete. Proceed to the NI U-Boot upgrade. If you are updating a second CMM (CMM B), go to Step 3.
3. Execute the **update uboot cmm slot** command to update the U-Boot File on Secondary CMM (CMM B). The command below is used if the Secondary CMM is in Slot 2 ("cmm 2"). If the Secondary CMM is in Slot 1, you would enter "cmm 1".

```
OS10K-> update uboot cmm 2 file u-boot.7.2.1.R02.266.tar.gz
```

WARNING: DO NOT INTERRUPT the upgrade process until it is complete ("Update success¬fully completed"). Interruption of the process will result in an unrecoverable failure condition.

NI Upgrade

Follow the steps below to upgrade the U-Boot File on the NI(s):

1. Execute the following CLI command to update the U-Boot File on NI(s).

```
OS10K-> update uboot ni all file u-boot.7.2.1.R02.266.tar.gz
```

2. When the "Update successfully completed" message appears, execute the following CLI command to delete the U-Boot File from the /flash directory:

```
OS10K-> rm u-boot.7.2.1.R02.266.tar.gz
```

IMPORTANT NOTE: Depending on the version of the 7.X build you are upgrading from, you may receive an error message when you execute the "**update uboot ni**" command. Simply re-enter the command, and the upgrade will proceed normally (shown below).

```
OS10K-> update uboot ni all file u-boot.7.2.1.R02.266.tar.gz
ERROR: Update failed for slot(s) 1 2 3 4 5 7
OS10K-> update uboot ni all file u-boot.7.2.1.R02.266.tar.gz
Please wait....Update successfully completed
```

4. Upgrade the image files

Follow the steps below to upgrade the image files to 7.3.2.R01:

Reload the switch from the working directory.

```
OS10K-> reload from working no rollback-timeout
```

After the switch finishes rebooting, log into the switch and copy the image files from the Working directory to the Certified directory.

If you have a single CMM switch enter:

```
OS10K-> copy running certified
```

If you have redundant CMMs enter:

```
OS10K-> copy running certified flash-synchro
```

5. Verify the Software Upgrade

To verify that the software was successfully upgraded to 7.3.2.R01, use the **show microcode** command as shown below.

```
OS10K-> show microcode
   Package      Release      Size          Description
---------+----------------+-------------+-------------------
Ros.img      7.3.2.344.R01    67784336     Alcatel-Lucent OS
Reni.img     7.3.2.344.R01    59189856     Alcatel-Lucent NI
```

OmniSwitch 6900 – Upgrade Instructions

Upgrading OS6900 Switches to 7.3.2.R01 consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support Website and download and unzip the 7.3.2.R01 upgrade files for the OS6900. The Zip File contains the following files:

- Image Files **-** Tos.img (7.3.2.R01)
- U-Boot File - u-boot.7.2.1.R02.266.tar.gz (if required)
- FPGA File - tor_fpgas_130_120.vme (if required)

2. FTP the Upgrade Files to the Switch

FTP the upgrade files to the following directories of the switch you are upgrading:

- U-Boot File - u-boot.7.2.1.R02.266.tar.gz - /**flash** directory (if required)
- FPGA File - tor_fpgas_130_120.vme - /**flash** directory (if required)
- Image File - Tos.img - **/flash/working** directory

3. Upgrading the U-Boot File (if required)

Follow the steps below to upgrade the U-Boot File:

Execute the following CLI command to update the U-Boot File on the switch.

```
OS6900-> update uboot cmm 1 file u-boot.7.2.1.R02.266.tar.gz
Sample output for "update uboot cmm 1”
u-boot.bin
u-boot.bin.md5sum
u-boot.bin: OK
Erasing blocks: 4/4 (100%)lease wait.
Writing data: 0k/512k (100%)
Verifying data: 0k/512k (100%)
U-boot successfully updated
Update successfully completed
```

WARNING: DO NOT INTERRUPT the upgrade process until it is complete (“Update success¬fully completed”). Interruption of the process will result in an unrecoverable failure condition.

4. Upgrading the FPGA (if required)

Follow the steps below to upgrade the FPGA File:

1. Execute the following CLI command to update the FPGA File on the switch.

```
OS6900-> update fpga cmm 1 file tor_fpgas_130_120.vme
Sample output for "update fpga cmm 1″
Wed Feb  8 11:27:59 : ChassisSupervisor MipMgr info message:
+++ Starting CMM FPGA Upgrade
OS6900 system and expansion fpga update
Please wait.........Update successfully completed
```

2. After the FPGA upgrade has successfully completed (“Update successfully completed”), delete the U- Boot and the FPGA Files from the /flash directory by entering the following CLI commands:

```
OS6900-> rm u-boot.7.2.1.R02.266.tar.gz
OS6900-> rm tor_fpgas_130_120.vme
```

5. Upgrade the image file

Follow the steps below to upgrade the image file:

Reload the switch from the working directory.

OS6900-> reload from working no rollback-timeout

After the switch finishes rebooting, log into the switch.

Copy the image files from the Working Directory to the Certified Directory by entering the following command:

```
OS6900-> copy running certified
```

6. Verify the Software Upgrade

To verify that the software was successfully upgraded to 7.3.2.R01, use the **show microcode** command as shown below:

```
OS6900-> show microcode
Package         Release          Size        Description
--------------+-----------------+------------+-------------------
Tos.img       7.3.2.344.R01     106031376     Alcatel-Lucent OS
```

Appendix B: Previous Release Feature Summary

Existing Hardware - AOS 7.1.1.R01

The following hardware was introduced with AOS Release 7.1.1.R01 for the OmniSwitch 10K.

OmniSwitch 10K Chassis

The OmniSwitch 10K is a high performance chassis accomodating high-density Gigabit Ethernet and 10-Gigabit Ethernet Network Interface (NI) modules.

- 8 Slots – Network Interface Modules
- 2 Slots – Chassis Management Modules (Integrated Management and Chassis Fabric Module)
- 2 Slots – Chassis Fabric Modules
- 2 Slots – Fan Trays (Two fan trays required)
- 4 Slots – Power Supplies

OS10K-CMM

The Chassis Management Module (CMM) provides both management and switching fabric for the OmniSwitch chassis. The CMM provides key system services and backup system services when a secondary CMM is present.

OS10K-CFM

The Chassis Fabric Module (CFM) provides the switching fabric for the chassis. Additional CFMs provide increased switching throughput, as well as redundancy.

OS10K-GNI-C48E

Provides 48 wire-rate RJ-45 1000Base-T ports and large table support for L2, L3, and ACL policies.

OS10K-GNI-U48E

Provides 48 wire-rate 1000BaseX SFP ports and large table support for L2, L3, and ACL policies.

OS10K-XNI-U32S

Provides 32 10-Gigabit SFP+ ports as well as support for 1-Gigabit SFP transceivers. Supports standard tables for L2, L3 and ACL policies.

OS10K-PS-25A

AC power supply auto-ranging from 110VAC-240VAC providing 1250W at 110VAC and 2500W at 240VAC.

OS10K-PS-24D

DC power supply providing up to 2400 watts of power with 36-72VDC input.

OS10K-Fan-Tray

Contains 12 individual variable-speed fans per tray.

Existing Hardware - AOS 7.2.1.R01

The following hardware was introduced in AOS Release 7.2.1.R01 for the OmniSwitch 6900.

OmniSwitch 6900-X20

- 10-Gigabit Ethernet fixed configuration chassis in a 1U form factor with 20 SFP+ ports, one optional module slot, redundant AC or DC power and front-to-rear cooling. The switch includes:
- 1 – Console Port (USB Form Factor - RS-232)
- 1 – USB Port (For use with Alcatel-Lucent **OS-USB-FLASHDR** USB flash drive)
- 1 – EMP Port
- 20 – SFP+ Ports
- 1 Slot– Optional module
- 1 Slot – Fan Tray
- 2 Slots – Power Supplies (AC or DC)

OmniSwitch 6900-X40

- 10-Gigabit Ethernet fixed conguration chassis in a 1U form factor with 40 SFP+ ports, two optional module slots, redundant AC or DC power and front-to-rear cooling. The switch includes:
- 1 – Console Port (USB Form Factor - RS-232)
- 1 – USB Port (For use with Alcatel-Lucent **OS-USB-FLASHDR** USB flash drive)
- 1 – EMP Port
- 40 – SFP+ Ports
- 2 Slots– Optional Modules
- 1 Slot – Fan Tray
- 2 Slots – Power Supplies (AC or DC)

OS-XNI-U4

10-Gigabit Ethernet module for the OS6900 series of swicthes with 4 SFP+ ports that support 1-Gigabit and 10-Gigabit transceivers.

OS-XNI-U12

10-Gigabit Ethernet module for the OS6900 series of swicthes with 12 SFP+ ports that support 1-Gigabit and 10-Gigabit transceivers.

OS6900-BP-F (YM-2451C) Power Supply

450W modular AC power supply with front-to-rear cooling.

OS6900-BPD-F (YM-2451D) Power Supply

450W modular DC power supply with front-to-rear cooling.

OS6900-FT-F FanTray

Contains 4 individual variable-speed fans per tray with front-to-rear cooling.

Existing Hardware - AOS 7.2.1.R02

The following hardware was introduced in AOS Release 7.2.1.R02.

NOTE: The hardware described below requires the GA build 7.2.1.323.R02.

OmniSwitch 6900 Rear-to-Front Cooling

The OmniSwitch 6900 now supports a rear-to-front cooling option with the rear-to-front fantray and power supply combination. Note the following:

- The airflow direction of the power supplies and fantray must be the same.
- The switch must be upgraded to the latest UBoot version 7.2.1.266.R02 to support rear-to-front cooling.

OS-QNI-U3 Module

40-Gigabit Ethernet module for the OS6900 series of swicthes with 3 QSFP+ ports that support 40-Gigabit transceivers. Note: Refer to the hot-swap section for hot-swap and module insertion requirements.

OS-HNI-U6 Module

10/40-Gigabit Ethernet module for the OS6900 series of swicthes with 4 SFP+ ports that support 1-Gigabit and 10-Gigabit transceivers and 2 QSFP+ ports that support 40-Gigabit transceivers. Note: Refer to the hot-swap section for hot-swap and module insertion requirements.

QSFP-40G-SR Transceiver

Four channel 40 Gigabit optical transceiver (QSFP+). Supports link lengths of 100m and 150m respectively on OM3 and OM4 multimode fiber cables. Note: Supports the required DDM parameters of Voltage (V) and Temperature (T) only.

QSFP-40G-C Transceiver

40-Gigabit direct attached copper cable available in 1/3/7 meter lengths.

OS6900-BP-R (YM-2451F) Power Supply

450W modular AC power supply with rear-to-front cooling.
Note: This power supply is differentiated from the front-to-rear power supplies by purple coloring.

OS6900-BPD-R (YM-2451P) Power Supply

450W modular DC power supply with rear-to-front cooling.
Note: This power supply is differentiated from the front-to-rear power supplies by purple coloring.

OS6900-FT-R FanTray

Contains 4 individual variable-speed fans per tray with rear-to-front cooling.
Note: This fan tray is differentiated from the front-to-rear fan tray by an **R->F** label and purple coloring.

Existing Hardware - AOS 7.3.1.R01

The following hardware was introduced in AOS Release 7.3.1.R01.

OS10K-XNI-U16L

OS10K network interface card includes 8 unpopulated 10G SFP+ ports (1-8) and 8 unpopulated 1G SFP+ ports (9-16). 1G ports can be updated to 10G through license upgrade. Supports standard tables for L2, L3 and ACL policies, MPLS ready.

OS10K-XNI-U16E
OS10K network interface card includes 16 unpopulated 10G SFP+ ports. Supports standard tables for L2, L3 and ACL policies, MPLS ready.

OS10K-XNI-U32E
OS10K network interface card includes 32 unpopulated 10G SFP+ ports. Supports standard tables for L2, L3 and ACL policies, MPLS ready.

OS10K-QNI-U4E
OS10K network interface card includes 4 unpopulated 40G QSFP+ ports. Supports standard tables for L2, L3 and ACL policies, MPLS ready.

OS10K-QNI-U8E
OS10K network interface card includes 8 unpopulated 40G QSFP+ ports. Supports standard tables for L2, L3 and ACL policies, MPLS ready.

QSFP-40G-LR Transceiver
Four channel 40 Gigabit optical transceiver (QSFP+). Supports single mode fiber over 1310nm wavelength with a typical reach 10 km. Note: Supports the DDM parameters of Voltage (V), Temperature (T), Current (mA) and Input (dBm). If the threshold values of the transceiver are '0' then NS (Not supported) will be displayed in the DDM output display.

SFP-10G-24DWD80 Transceiver
10 Gigabit DWDM optical tranceiver with an LC connector. Supports single mode fiber over 1558.17nm with a typical reach of 80 km. Note: Only supported on XNI (10G) modules.

SFP-10G-GIG-SR Transceiver

Dual-speed SFP+ optical transceiver. Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Supports 1000BaseSX and 10GBASE-SR.

Existing Software Features Summary - AOS 7.1.1.R01

The following software features were introduced in the 7.1.1.R01 release for the OmniSwitch 10K, subject to the feature exceptions and problem reports described in the 7.1.1.R01 Release Notes:

AOS 7.1.1. R01 Feature Summary Table

Feature	Platform	Software Package
Manageability Feature Support		
CLI	OS10K	Base
Ethernet Interfaces	OS10K	Base
ISSU	OS10K	Base
Multiple VRF Routing and Forwarding	OS10K	Base
Network Time Protocol (NTP)	OS10K	Base
Pause Control/Flow Control	OS10K	Base
Remote Access FTP SCP SSH/SFTP Telnet TFTP	OS10K	Base
Smart Continuous Switching Hot Swap Management Module Failover Power Monitoring Redundancy	OS10K	Base
SNMP	OS10K	Base
Software Rollback – Multi-Image/Multi-Config	OS10K	Base
Storm Control	OS10K	Base
Text File Configuration	OS10K	Base
UDLD	OS10K	Base
USB Support	OS10K	Base
Web-Based Management (WebView)	OS10K	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS10K	Base
802.1Q	OS10K	Base
Configurable Hash Mode	OS10K	Base

Feature	Platform	Software Package
Link Aggregation –Static and LACP (802.3ad)	OS10K	Base
Multi-Chassis Link Aggregation	OS10K	Base
Source Learning	OS10K	Base
Spanning Tree • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard	OS10K	Base
VLANs	OS10K	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS10K	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS10K	Base
BGP4 with Graceful Restart	OS10K	Base
DNS Client	OS10K	Base
GRE	OS10K	Base
IP Multicast Routing	OS10K	Base
IP Multicast Switching (IGMP)	OS10K	Base
IP Multicast Switching (Proxying)	OS10K	Base
IP Multinetting	OS10K	Base
IP Route Map Redistribution	OS10K	Base
IP-IP Tunneling	OS10K	Base
OSPFv2	OS10K	Base
RIPv1/v2	OS10K	Base
Routing Protocol Preference	OS10K	Base
Server Load Balancing	OS10K	Base
VRRPv2	OS10K	Base

Feature	Platform	Software Package
IPv6 Feature Support		
BGP4 BGP IPv6 Extensions	OS10K	Base
IPSec IPv6 OSPFv3 RIPng	OS10K	Base
IPv6 Client and/or Server Support	OS10K	Base
IPv6 Multicast Routing	OS10K	Base
IPv6 Multicast Switching (MLD v1/v2)	OS10K	Base
IPv6 Routing	OS10K	Base
IPv6 Scoped Multicast Addresses	OS10K	Base
IPv6 Neighbor Discovery Support	OS10K	Base
OSPFv3	OS10K	Base
RIPng	OS10K	Base
VRRPv3	OS10K	Base
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS10K	Base
Ingress and egress bandwith shaping	OS10K	Base
Policy Based Routing	OS10K	Base
Tri-Color Marking	OS10K	Base
Multicast Feature Support		
DVMRP	OS10K	Base
IGMP Multicast Group Configuration Limit	OS10K	Base
IGMP Relay	OS10K	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS10K	Base
L2 Static Multicast Address	OS10K	Base
PIM / PIM-SSM (Source-Specific Multicast)	OS10K	Base
Monitoring/Troubleshooting Feature Support		

Feature	Platform	Software Package
DDM - Digital Diagnostic Monitoring	OS10K	Base
Health Statistics	OS10K	Base
Ping and Traceroute	OS10K	Base
Policy Based Mirroring	OS10K	Base
Port Mirroring	OS10K	Base
Port Monitoring	OS10K	Base
Remote Port Mirroring	OS10K	Base
Rmon	OS10K	Base
sFlow	OS10K	Base
Switch Logging and Syslog	OS10K	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS10K	Base
Ethernet Services	OS10K	Base
L2 Control Protocol Tunneling (L2CP)	OS10K	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS10K	Base
Account & Password Policies	OS10K	Base
Admin User Remote Access Restriction Control	OS10K	Base
ARP Defense Optimization	OS10K	Base
ARP Poisoning Detect	OS10K	Base
Authenticated Switch Access	OS10K	Base
IP DoS Filtering	OS10K	Base
Learned Port Security (LPS)	OS10K	Base
Policy Server Management	OS10K	Base

Existing Software Features Summary – AOS 7.2.1.R01

The following software features were introduced in the 7.2.1.R01 release for the OmniSwitch 6900, subject to the feature exceptions and problem reports described later in the 7.2.1.R01 Release Notes:

Features listed as 'Base' are included as part of the base software and do not require any license installation. Features listed as 'Advanced' require the installation of an Advanced license.

AOS 7.2.1.R01 Feature Summary Table

Feature	Platform	License
Manageability Feature Support		
CLI	OS6900	Base
Ethernet Interfaces	OS6900	Base
License Management	OS6900	Base
Multiple VRF Routing and Forwarding	OS6900	Advanced
Network Time Protocol (NTP)	OS6900	Base
Pause Control(RX) /Flow Control	OS6900	Base
Remote Access FTP SCP SSH/SFTP Telnet TFTP	OS6900	Base
Resiliency Features Hot Swap Expansion Modules Power Supply Redundancy Fan Redundancy	OS6900	Base
SNMP	OS6900	Base
Software Rollback – Multi-Image/Multi-Config	OS6900	Base
Storm Control	OS6900	Base
Text File Configuration	OS6900	Base
UDLD	OS6900	Base
USB Support	OS6900	Base
Web-Based Management (WebView)	OS6900	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS6900	Base
802.1Q	OS6900	Base

Feature	Platform	License
Configurable Hash Mode	OS6900	Base
HA-VLAN	OS6900	Base
Link Aggregation –Static and LACP (802.3ad)	OS6900	Base
Multi-Chassis Link Aggregation	OS6900	Base
MVRP	OS6900	Base
Source Learning	OS6900	Base
Spanning Tree • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard	OS6900	Base
Universal Network Profiles (UNP)	OS6900	Base
VLANs	OS6900	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS6900	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS6900	Base
BGP4 with Graceful Restart	OS6900	Advanced
DNS Client	OS6900	Base
GRE	OS6900	Base
IP Multicast Routing	OS6900	Advanced
IP Multicast Switching (IGMP)	OS6900	Base
IP Multicast Switching (Proxying)	OS6900	Base
IP Multinetting	OS6900	Base
IP Route Map Redistribution	OS6900	Base
IP-IP Tunneling	OS6900	Base
OSPFv2	OS6900	Advanced
RIPv1/v2	OS6900	Base
Routing Protocol Preference	OS6900	Base

Feature	Platform	License
Server Load Balancing	OS6900	Base
VRRPv2	OS6900	Advanced
IPv6 Feature Support		
BGP4 BGP IPv6 Extensions	OS6900	Advanced
IPSec IPv6 OSPFv3 RIPng	OS6900	Advanced
IPv6 Client and/or Server Support	OS6900	Base
IPv6 Multicast Routing	OS6900	Advanced
IPv6 Multicast Switching (MLD v1/v2)	OS6900	Base
IPv6 Routing	OS6900	Advanced
IPv6 Scoped Multicast Addresses	OS6900	Base
IPv6 Neighbor Discovery Support	OS6900	Base
OSPFv3	OS6900	Advanced
RIPng	OS6900	Advanced
VRRPv3	OS6900	Advanced
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS6900	Base
Ingress and egress bandwith shaping	OS6900	Base
Policy Based Routing	OS6900	Advanced
Tri-Color Marking	OS6900	Base
Multicast Feature Support		
DVMRP	OS6900	Advanced
IGMP Multicast Group Configuration Limit	OS6900	Base
IGMP Relay	OS6900	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS6900	Base
L2 Static Multicast Address	OS6900	Base
PIM / PIM-SSM (Source-Specific	OS6900	Advanced

Feature	Platform	License
Multicast)		
Monitoring/Troubleshooting Feature Support		
DDM - Digital Diagnostic Monitoring	OS6900	Base
Health Statistics	OS6900	Base
Ping and Traceroute	OS6900	Base
Policy Based Mirroring	OS6900	Base
Port Mirroring	OS6900	Base
Port Monitoring	OS6900	Base
Remote Port Mirroring	OS6900	Base
Rmon	OS6900	Base
sFlow	OS6900	Base
Switch Logging and Syslog	OS6900	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS6900	Base
Ethernet Services	OS6900	Base
L2 Control Protocol Tunneling (L2CP)	OS6900	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS6900	Base
Account & Password Policies	OS6900	Base
Admin User Remote Access Restriction Control	OS6900	Base
ARP Defense Optimization	OS6900	Base
ARP Poisoning Detect	OS6900	Base
Authenticated Switch Access	OS6900	Base
IP DoS Filtering	OS6900	Base
Learned Port Security (LPS)	OS6900	Base
Policy Server Management	OS6900	Base

Existing Software Features Summary – AOS 7.2.1.R02

The following software features were introduced in the 7.2.1.R02 release, subject to the feature exceptions and problem reports described later in the 7.2.1.R02 Release Notes:

Features listed as 'Base' are included as part of the base software and do not require any license installation. Features listed as 'Advanced' require the installation of an Advanced license.

Feature	Platform	License
Layer 2 Feature Support		
High Availability VLAN		
• Added support for OS10K	OS10K	Base
• HA-VLAN with MCLAG	OS10K/6900	Base
Multi-Chassis Link Aggregation		
• Configurable Chassis Group ID (Multiple MC-LAG Domains)	OS10K/6900	Base
• Standalone Port in VIP VLAN	OS10K/6900	Base
• SLB Over MC-LAG	OS10K/6900	Base
MVRP		
• Added support for OS10K	OS10K	Base
Universal Network Profiles		
• UNP with Dynamic Profiles	OS6900	Base
• UNP with Link-Aggregation	OS6900	Base
• UNP with MC-LAG	OS6900	Base
• UNP with Learned Port Security	OS6900	Base
Layer 3 Feature Support		
16 ECMP routes for IPv6	OS10K/6900	Base
Qos		
VFC/VoQ Profiles		
• Added support for profiles 2-4	OS10K/6900	Base
• Added support for WRED	OS6900	Base
Security		
Learned Port Security Enhancements	OS10K/6900	Base

Existing Software Features Summary – AOS 7.3.1.R01

The following software features were introduced in the 7.3.1.R01 release for the OmniSwitch, subject to the feature exceptions and problem reports described later in the 7.3.1.R01 Release Notes:

Features listed as 'Base' are included as part of the base software and do not require any license installation. Features listed as 'Advanced' require the installation of an Advanced license.

Feature	Platform	License
Data Center Feature Support		
Shortest Path Bridging (SPB)	OS10K/6900	Advanced
Data Center Bridging • DCBX • ETS • PFC	 OS10K/6900 OS10K/6900 OS10K/6900	 Data Center Data Center Data Center
Edge Virtual Briding (EVB)	OS10K/6900	Data Center
Virtual Network Profiles • SAP/SPB-M Services • Customer Domains (Multi-tenancy) • Dynamic SAP • UNP over MC-LAG on OS10K	 OS10K/6900 OS10K/6900 OS10K/6900 OS10K/6900	 Base Base Base Base
Layer 2 Feature Support		
Ethernet Ring Protection v2 (ERPv2)	OS10K/6900	Base
Layer 3 Feature Support		
VRF Management	OS10K/6900	Base
VRF Route Leak	OS10K/6900	Base
Management Feature Support		
Virtual Chassis	OS10K/6900	Advanced
SFP+ Line Diags & Enhanced Port Performance (EPP)	OS10K/6900	Base
License Management	OS10K/6900	Base
Ethernet OAM • ITU Y1731 and 802.1ag	OS10K/6900 OS10K/6900	Base

Feature	Platform	License
Service Assurance Agent	OS10K/6900	Base

Note: The SAP/SPB-M Services, Customer Domains, Dynamic SAP, and Virtual Chassis features were introduced in AOS Release 7.3.1.632.R01. The remaining features in this section were introduced in AOS Release 7.3.1.519.R01.