# Release Notes - Rev. A

## OmniSwitch
## 6900/6860(E)/6865/6560/9900

### Release 8.4.1.R02

These release notes accompany release 8.4.1.R02. 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.

## Contents

## Related Documentation

These release notes should be used in conjunction with OmniSwitch AOS Release 8 User Guides. The following are the titles of the user guides that apply to this release. User guides can be downloaded at: https://support.esd.alcatel-lucent.com

- OmniSwitch 6900 Hardware User Guide
- OmniSwitch 6860(E) Hardware User Guide
- OmniSwitch 6865 Hardware User Guide
- OmniSwitch 6560 Hardware User Guide
- OmniSwitch 9900 Hardware User Guide
- OmniSwitch AOS Release 8 CLI Reference Guide
- OmniSwitch AOS Release 8 Network Configuration Guide
- OmniSwitch AOS Release 8 Switch Management Guide
- OmniSwitch AOS Release 8 Advanced Routing Configuration Guide
- OmniSwitch AOS Release 8 Data Center Switching Guide
- OmniSwitch AOS Release 8 Specifications Guide
- OmniSwitch AOS Release 8 Transceivers Guide

## System Requirements

### Memory Requirements

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

| Platform | SDRAM | Flash |
|---|---|---|
| OS6900-X Models | 2GB | 2GB |
| OS6900-T Models | 4GB | 2GB |
| OS6900-Q32 | 8GB | 2GB |
| OS6900-X72 | 8GB | 4GB |
| OS6860(E) | 2GB | 2GB |
| OS6865 | 2GB | 2GB |
| OS6560 | 2GB | 2GB |
| OS9900 | 16GB | 2GB |

### 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. Use the '**show hardware-info**' command to determine the current versions.

Switches not running the minimum version required should upgrade to the latest UBoot or FPGA that is available with the 8.4.1.R02 AOS software 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.

### OmniSwitch 6900-X20/X40 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum UBoot | Minimum FPGA |
|---|---|---|
| CMM (if XNI-U12E support is not needed) | 7.2.1.266.R02 | 1.3.0/1.2.0 |
| CMM (if XNI-U12E support is needed) | 7.2.1.266.R02 | 1.3.0/2.2.0 |
| All Expansion Modules | N/A | N/A |

### OmniSwitch 6900-T20/T40 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum UBoot | Minimum FPGA |
|---|---|---|
| CMM (if XNI-U12E support is not needed) | 7.3.2.134.R01 | 1.4.0/0.0.0 |
| CMM (if XNI-U12E support is needed) | 7.3.2.134.R01 | 1.6.0/0.0.0 |
| All Expansion Modules | N/A | N/A |

### OmniSwitch 6900-Q32 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum UBoot | Minimum FPGA |
|---|---|---|

| Hardware | Minimum UBoot | Minimum FPGA |
|---|---|---|
| CMM<br>All Expansion Modules | 7.3.4.277.R01<br>N/A | 0.1.8<br>N/A |

### OmniSwitch 6900-X72 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum Uboot | Minimum FPGA |
|---|---|---|
| CMM<br>All Expansion Modules | 7.3.4.31.R02<br>N/A | 0.1.10<br>N/A |

### OmniSwitch 6860(E) – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum Uboot | Minimum FPGA* |
|---|---|---|
| OS6860/OS6860E (except U28) | 8.1.1.70.R01 | 0.9 (0x9) |
| OS6860E-U28 | 8.1.1.70.R01 | 0.20 (0x14) |
| OS686E-P24Z8 | 8.4.1.17.R01 | 0.5 (0x5) |

***Note**: In previous AOS releases the FPGA version was displayed in hexadecimal format. Beginning in 8.4.1.R01 it is displayed in decimal format.

### OmniSwitch 6865 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum Uboot | Minimum FPGA* |
|---|---|---|
| OS6865-P16X | 8.3.1.125.R01 | 0.20 (0x14) (minimum)<br>0.22 (0x16) (current) |
| OS6865-U12X | 8.4.1.17.R01 | 0.23 (0x17) |
| OS6865-U28X | 8.4.1.17.R01 | 0.11 (0xB) |

***Note**: In previous AOS releases the FPGA version was displayed in hexadecimal format. Beginning in 8.4.1.R01 it is displayed in decimal format.

### OmniSwitch 6560 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Minimum Uboot | Minimum FPGA |
|---|---|---|
| OS6560-P24Z24 | 8.4.1.23.R02 | 0.6 (0x6) |
| OS6560-P24Z8 | 8.4.1.23.R02 | 0.6 (0x6) |

### OmniSwitch 9900 – AOS Release 8.4.1.229.R02(GA)

| Hardware | Coreboot-uboot | Control FPGA | Power FPGA |
|---|---|---|---|
| OS99-CMM | 8.3.1.103.R01 | 2.3.0 | 0.8 |
| OS9907-CFM | 8.3.1.103.R01 | - | - |
| OS99-GNI-48 | 8.3.1.103.R01 | 1.2.4 | 0.9 |
| OS99-GNI-P48 | 8.3.1.103.R01 | 1.2.4 | 0.9 |
| OS99-XNI-48 | 8.3.1.103.R01 | 1.3.0 | 0.6 |
| OS99-XNI-U48 | 8.3.1.103.R01 | 2.9.0 | 0.8 |
| OS99-GNI-U48 | 8.4.1.166.R01 | 0.3.0 | 0.2 |

## [IMPORTANT] *MUST READ*: AOS Release 8.4.1.R02 Prerequisites and Deployment Information

### General Information

- ***Note: Early availability features are available in AOS and can be configured. However, they have not gone through the complete AOS validation cycle and are therefore not officially supported.***
- Please refer to the Feature Matrix in Appendix A for detailed information on supported features for each platform.
- Prior to upgrading to AOS Release 8.4.1.R02 please refer to Appendix B for important best practices, prerequisites, and step-by-step instructions.
- Many of the "**show ip multicast**" commands have changed to remove the "**domain**" filter keyword which was introduced in 8.4.1.R01. Please refer to the CLI Reference Guide for additional details on the IPMS CLI changes.

### Additional Information

- All switches that ship from the factory will default to VC mode (requiring a vcboot.cfg configuring file) and attempt to run the automatic VC, automatic remote configuration, and automatic fabric protocols. Please note that since the switches default to VC mode, automatic remote configuration does not support the downloading of a 'boot.cfg' file, only the 'vcboot.cfg' file is supported.

**Note**: None of the ports on the OS6865 models default to auto-vfl so automatic VC will not run by default on newly shipped switches. However, automatic remote configuration and automatic fabric will run by default.

- Beginning in 8.4.1.R01 when configuring BFD with protocols that support echo-only mode (VRRP or static routes) the configuration of a Loopback0 address is required. Upgrading from a previous release without a Loopback0 address will result in a configuration error.

## Licensed Features

The table below lists the licensed features in this release and whether or not a license is required for the various models.

| | Data Center License Installation Required? | | | | |
|---|---|---|---|---|---|
| | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 |
| Data Center Features | | | | | |
| | | | | | |
| DCB (PFC,ETS,DCBx) | Yes | N/S | N/S | N/S | N/S |
| EVB | Yes | N/S | N/S | N/S | N/S |
| FIP Snooping | Yes | N/S | N/S | N/S | N/S |
| FCoE VXLAN | Yes | N/S | N/S | N/S | N/S |

## CodeGuardian

Alcatel-Lucent Enterprise and LGS Innovations have combined to provide the first network equipment to be hardened by an independent group. CodeGuardian promotes security and assurance at the network device level using independent verification and validation of source code, software diversification to prevent exploitation and secure delivery of software to customers.

CodeGuardian employs multiple techniques to identify vulnerabilities such as software architecture reviews, source code analysis (using both manual techniques and automated tools), vulnerability scanning tools and techniques, as well as analysis of known vulnerabilities in third party code.

### Software diversification

Software diversification randomizes the executable program so that various instances of the same software, while functionally identical, are arranged differently. The CodeGuardian solution rearranges internal software while maintaining the same functionality and performance and modifies the deliverable application to limit or prevent/impede software exploitation. There will be up to 3 different diversified versions per GA release of code.

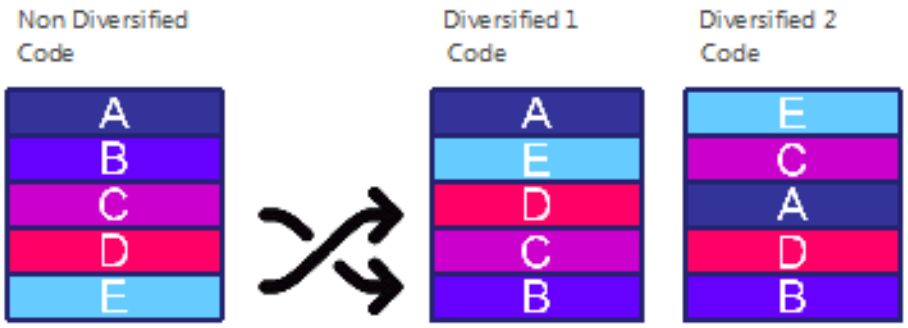


### CodeGuardian AOS Releases

| Standard AOS Releases | AOS CodeGuardian Release | LGS AOS CodeGuardian Release |
|---|---|---|
| AOS 8.4.1.R02 | AOS 8.4.1.RX2 | AOS 8.4.1.LX2 |

- X=Diversified image 1-3
- ALE will have 3 different diversified images per AOS release (R12 through R32)
- Our partner LGS will have 3 different diversified images per AOS release (L12 through L32)

Please contact customer support for additional information.

## New / Updated Hardware Support

The following new hardware is being introduced in this release.

**OS6560-P24Z24**

Fixed configuration chassis in a 1U form factor with:

- Twenty-four (24) - 100/1000/2.5G Base-T 802.3bt PoE ports (95W HPoE)
- Four (4) - SFP+ (1G/10G) ports
- Two (2)- 20G virtual chassis VFL ports
- USB port
- RJ-45 console port

**OS6560-P24Z8**

Fixed configuration chassis in a 1U form factor with:

- Sixteen (16) - 10/100/1000 BaseT 802.3at PoE ports (30W PoE)
- Eight (8) - 100/1000/2.5G Base-T 802.3bt PoE ports (95W HPoE)
- Two (2) - SFP+ (1G/10G) ports
- USB port
- RJ-45 console port

**OS6560-P48Z16 (Future Availability)**

Fixed configuration chassis in a 1U form factor with:

- Sixteen (32) - 10/100/1000 BaseT 802.3at PoE ports (30W PoE)
- Sixteen (16) - 100/1000/2.5G Base-T 802.3bt PoE ports (95W HPoE)
- Four (4) - SFP+ (1G/10G) ports
- Two (2) – 20G virtual chassis VFL ports
- USB port
- RJ-45 console port

**OS6560-BP-P**

300W AC power supply providing both system and PoE power.

**OS6560-BP-PH**

600W AC power supply providing both system and PoE power.

**OS6560-BP-PX**

920W AC power supply providing both system and PoE power.

**Transceivers**

Support for the following transceivers has been added for the OS6560. Please refer to the Transceivers Guide for additional details on other existing transceivers and the supported platforms.

| 100M | Gigabit | Dual-Speed | 10G | 40G/VFL |
|---|---|---|---|---|
| SFP-100-BX-LC-D | SFP-GIG-SX | SFP-DUAL-MM | SFP-10G-SR | QSFP-40G-SR |
| SFP-100-BX-LC-U | SFP-GIG-LX | SFP-DUAL-MM-N | SFP-10G-LR | QSFP-40G-C1M/3M/40CM |
| SFP-100-LC-MM | SFP-GIG-LH40 | SFP-DUAL-SM10 | SFP-10G-ER | - OS6560-CBL-100 |
| SFP-100-LC-SM15 | SFP-GIG-LH70 | SFP-DUAL-BX-D | SFP-10G-ZR | - OS6560-CBL-300 |
| SFP-100-LC-SM40 | SFP-GIG-T | SFP-DUAL-BX-U | SFP-10G-T | - OS6560-CBL-40 |
| | SFP-GIG-EXTND | | SFP-10G-C1M/3M/7M/60CM | QSFP-40G-AOC20M |
| | SFP-GIG-BX-D | | SFP-10G-GIG-SR | |
| | SFP-GIG-BX-U | | SFP-10G-GIG-LR | |
| | SFP-GIG-BX-D20 | | | |
| | SFP-GIG-BX-U20 | | | |
| | SFP-GIG-BX-D40 | | | |
| | SFP-GIG-BX-U40 | | | |

## New Software Features and Enhancements

The following software features are being introduced with the 8.4.1.R02 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 "Data Center" require a license to be installed.

### 8.4.1.R02 New Feature/Enhancements Summary

| Feature | Platform |
|---|---|
| OS9900 - VC of 2 | OS9900 |
| Access Guardian Integration with UPAM for Guest Access and BYOD | OS6860/OS6865/OS6560 |
| AG enhancement to support local time and location based policy in the presence of redirect server | All |
| Guest Tunneling Endpoint Switch<br><br>Guest Tunnel Termination Switch (GTTS) | OS6860/OS6865<br><br>OS6860/OS6865/6900-Q32/X72 |
| mDNS/SSDP (UPnP/DLNA) Relay Across L3 Networks | OS6860/OS6865/OS6900/OS9900 |
| WebView OmniAccess Stellar AP Management | OS6860/OS6865/OS6560/OS9900 |
| LLDP AOS Enhancements for Assignment of WLAN Parameters to the OmniAccess Stellar APs | OS6860/OS6865/OS6560/OS9900 |
| OmniAccess Stellar AP Integration with the OmniSwitch | All |
| AOS Enhancement for Dynamic WLAN VLAN Creation | OS6860/OS6865/OS6560/OS9900 |
| 1588-PTP on 6900-X72 | OS6900-X72 |
| MAC Authentication Enhancement | All |
| UNP CLI Enhancement | All |
| IPMS Profiles | OS6860/OS6865/OS6900 |

**Virtual Chassis for OS9900**
The OS9900 now supports a Virtual Chassis of 2.

- This release supports static VFLs only.
- It is recommended to use 10G VFL links (See PR 228215).
- Using a single-CMM configuration in OS9900 VC of 2 is a single point of failure that could lead to a full failure of the VC. It is highly recommended to use Dual-CMM chassis when operating a VC of 2. (See PR 228715)

**Access Guardian Integration with UPAM for Guest Access and BYOD**
In previous releases Access Guardian worked with ClearPass for Guest Access and BYOD solutions. This enhancement integrates Access Guardian with the new Unified Policy Access Manager (UPAM) solution for Guest Access and BYOD.

**AG enhancement to support local time and location based policy in the presence of redirect server**
Allows time and location based configuration in a UNP profile to be done locally on the OmniSwitch even if redirect is enabled.

**Guest Traffic Segregation - Guest Tunneling Using UNP/Services**

Guest Tunneling is a mechanism that is used to identify and isolate guest traffic from the rest of the internal network traffic. The tunneling protocol used is Layer 2 Generic Routing Encapsulation (GRE). A GRE tunnel is defined by configuring one end of the tunnel on an edge (access) switch and the other end of the tunnel on a Guest Tunnel Termination Switch (GTTS).

- Traffic received on the edge switch is classified into a UNP guest profile that is mapped to a Layer 2 GRE tunnel. The guest traffic is then encapsulated and tunneled through the network to the GTTS.
- When the tunneled traffic reaches the GTTS, the GRE encapsulation is removed and the traffic is then forwarded through a physical loopback port to a VLAN domain. At this point, the guest traffic can gain access to a perimeter network and/or the Internet.

**mDNS/SSDP (UPnP/DLNA) relay across L3 networks**

The Zero configuration for Multicast DNS (mDNS) and Simple Service Discovery Protocol (SSDP) is developed to extend mDNS and SSDP across Layer 3.

The zero configuration mDNS and SSDP solution allows:

- mDNS and SSDP compatible devices to discover network services across IP subnet boundaries.
- Selective sharing of network services based on sharing rules for mDNS capable devices. Sharing rules are defined based on VLAN, access role profile (UNP), and location.
- To provide the solution that is unified across wire or wireless (Aruba AP) network.
- Multicast optimization over the wireless (Aruba AP) network.

The mDNS or SSDP packet handling across layer 3 supports the following mode of operation:

- **Tunnel (Aruba) Mode**: Supports mDNS or SSDP compatible devices with Aruba controller with GRE tunnel protocol type 0x0. This is the default mode of operation.
- **Tunnel Standard Mode**: Supports tunneling for mDNS compatible devices to an OmniSwitch responder with GRE tunnel protocol type 0x6558. Only mDNS over IPv4 is supported.
- **Gateway Mode**: Supports mDNS or SSDP compatible devices to discover network services across IP subnet boundaries or VLANs. Only mDNS or SSDP over IPv4 is supported.
- **Responder Mode**: Supports mDNS compatible devices with OmniSwitch as a core switch (Responder). Only mDNS over IPv4 is supported.

**WebView OmniAccess Stellar AP Management**

The Cluster Virtual IP address to access the group of APs through OmniSwitch WebView can be automatically configured. The OmniSwitch acquires the Cluster Virtual IP address from the LLDP TLV received from the access points (AP).

All AP belonging to the same L2 domain and having the same cluster-ID are grouped into a single cluster. Each of these APs have their own unique IP address and the cluster is associated with a single virtual IP address for management.

The cluster can be configured or managed through a Web interface by connecting to the cluster virtual IP address. The cluster virtual IP address is associated with the primary AP of the cluster. The OmniSwitch automatically configures the cluster virtual IP address from the received LLDP packets from the APs.

**LLDP AOS Enhancements for Assignment of WLAN Parameters to the OmniAccess Stellar APs**

The OmniSwitch can advertise the WLAN management VLAN information and Access Point Location information of the APs connected to it using TLVs.

The WLAN Management VLAN is transmitted to AP through LLDP using existing Port VLAN TLV. The WLAN management VLAN is locally maintained for each port on the switch. The TLV must be enabled to advertise the information.

**OmniAccess Stellar AP Integration with the OmniSwitch**

Access Guardian provides the framework through which OmniAccess Stellar Access Points (APs) connected to an OmniSwitch are detected, learned, and managed. Wireless client traffic is then forwarded from the AP device to the OmniSwitch and onto the wired network. This integration provides a unified wireless over wired network access solution.

The OmniSwitch boots up with specific default configuration and operational settings that trigger the following process to detect, learn, and classify connected Stellar AP devices:

1. The switch and any Stellar AP device that is connected to an 802.1x port initially exchange Link Layer Detection Protocol (LLDP) TLV packets. Through this exchange of LLDP packets, the switch identifies and learns the device MAC address as an AP.
2. The detection of an AP device on an 802.1x port triggers the following actions that will automatically change the operational status of the specified options (the operational status overrides the configured status).
    - The transmission of LLDP Port VLAN ID and AP Location TLVs is operationally enabled on the UNP bridge port.
    - The trust tag status for the 802.1x port is operationally enabled.
    - The global status for dynamic VLAN configuration is operationally enabled for the switch.
3. Once the AP MAC address is detected and learned, a built-in LLDP UNP classification rule for access points classifies the AP device into a built-in default profile (defaultWLANProfile). The profile is associated with a VLAN into which the AP device is classified. This establishes a VLAN-port association (VPA) between the 802.1x port and profile VLAN on which the AP MAC address is learned and forwarded.
4. After the AP device connection is established, classified, and the management VLAN assigned, any of the following actions can occur:
    - The AP device sends DHCP packets.
    - The switch transmits LLDP packets to the AP device to advertise the management VLAN and AP location information.

- The AP device starts to send client-tagged traffic (tagged with the SSID VLAN). The switch will trust the VLAN tag of the AP client traffic and attempt to assign the traffic to a switch VLAN that matches the tag of the client traffic. If a matching switch VLAN does not exist, then the switch will dynamically create the necessary VLAN on which to forward the AP client traffic.

5. MVRP will then propagate the VLAN configuration (AP management VLAN and any static or dynamic VLAN that was automatically tagged to carry AP client traffic) to adjoining switches in the network. This process creates specific VLAN domains through which the untagged AP management traffic and tagged wireless client traffic is forwarded on the wired network.

The OmniSwitch detection and integration of OmniAccess Stellar APs results in a switch configuration that includes a management VLAN for the AP device and additional VLANs for wireless client-tagged traffic that is forwarded by the AP onto the wired network.

**AOS Enhancement for Dynamic WLAN VLAN Creation**
When an Access Point is detected, this enhancement supports the creation of dynamic WLAN VLAN from the tagged client packets on a port. However, to ensure the traffic to flow across the network, the created VLAN should get propagated through the network with the help of MVRP. Since MVRP supports only in flat mode, this solution works only in flat mode.

**IEEE 1588 – Precision Time Protocol (PTP)**
The Precision Time Protocol (PTP) is a protocol used to synchronize clocks throughout a computer network. On a local area network, it achieves clock accuracy in the sub-microsecond range, making it suitable for measurement and control systems. The intermediate switches must have the ability to support PTP and thereby update the link and/or residency time of frames in these switches - a concept knows as transparent clocking. There are two types of Transparent clocks, end-2-end Transparent Clock and peer-2-peer transparent clock. OmniSwitch supports one step End-2-End Transparent clock as defined in IEEE 1588 v2 Precision Time Protocol standard (supports both 1588v1/v2 Transparent Clock).

The “interfaces ptp admin-state” command can be used to enable or disable IEEE 1588 PTP time stamping on the switch, and set the internal priority for the incoming PTP packet.

**MAC Authentication Enhancement**

Added ‘unp 802.1x-pass-through’ command.

**UNP CLI Enhancement**

Added "0" support to the ‘aaa inactivity-logout’ command to avoid MAC address timeout.

**IPMS Profiles**

An IPMS profile is used to apply a pre-defined configuration to the global IPMS instance (all VLAN and service instances) or to a specific VLAN or service instance. Using a configuration profile to configure IPMS functionality avoids having to configure each IPMS parameter with a separate CLI command.

There is a “default” profile that defines a default set of IPMS parameter values that is automatically assigned to an IPMS instance. The default profile cannot be deleted, but the profile parameter values are configurable through this command.

## Open Problem Reports and Feature Exceptions

The problems listed here include problems known at the time of the product's release.

**General / System / Display**

| PR | Description | Workaround |
|---|---|---|
| 224454 | HAVLAN and distributed ARP cannot exist on same unit. | There is no known workaround at this time. |
| 224874 | On an OS6900-Q32 for broadcast and multicast traffic, "Input packets" in "show interface traffic" CLI and "InMcastPkts/ InBcastPkts" in "show interface counters" CLI is shown as "0" though traffic flows fine. Also "show interfaces accounting" CLI doesn't display total of TX and RX packets. | There is no known workaround at this time. This is a display issue only. |
| 225423 | With multiple users ingressing simultaneously on VXLAN port users are sometimes moved to "No VXLAN resource" and traffic drops for those users. | There is no known workaround at this time. |
| 226670 | On an OS6560 the "show lldp port remote-system" is not getting updated even after 30 seconds after bringing the port back up. | There is no known workaround at this time. |
| 226710 | '$' (duplicate static address) should not be shown in mac-table when mac becomes active on the port. | There is no known workaround at this time. There is no functional impact. |
| 227817 | The UDLD port status goes to "undetermined" state, instead of bidirectional after the second takeover. | After the second take-over, the port will come up within 360 seconds, then the UDLD functionality works correctly. |
| 228499 | Deleting linkagg which was part of DHL link displays the following "ERROR: Aggregate has DHL config". | Rebooting the switch will clear the error. |
| 228563 | MVRP auto-fabric may not start sometimes. | Use "auto-fabric discovery start" command to start manually. |
| 228722 | AOS is dropping the dhcp-discover from clients if AOS has a dhcp-client interface configured with an IP acquired from dhcp-server. | There is no known workaround at this time. |

**Hardware**

| PR | Description | Workaround |
|---|---|---|
| 228234 | The link doesn't come up when connecting 10G SFP+ CU 7M cable (120390-90; 2GSPWWX-NVG-EF) between two OS6865-U28X 10G ports. | Use a different 10G cable. |
| 224473 | 100M half-duplex detected when changing interface speed from auto to 100M or on multiple link toggle. | Toggle the interface. |

**OpenFlow**

| PR | Description | Workaround |
|---|---|---|
| 199279 | An interface which is a VPA of an Openflow VLAN dynamically learns MAC addresses on an OS6900. | There is no known workaround at this time. |
| 222968 | The traffic is not forwarded for all 224K MAC entries learned in hardware as OpenFlow L2-dest flows. Traffic forwarding is happening only for approximately 213K flows. | There is no known workaround at this time. |

**QoS**

| PR | Description | Workaround |
|---|---|---|
| 222202 | To apply a QoS config file with 511 entries on an OS9900 takes 10 minutes and results in a client unavailable error message on the console. | There is no known workaround at this time. |
| 225587 | On and OS9900 or OS6560 control packets are dropped due to any QoS policy which has disposition drop. | User needs to take care that control BPDUs are not qualified in any policy which might have policy action as disposition drop. |
| 227157 | For 9900 and 6560, ‘qos port maximum bandwidth’ behaves differently for ingress and egress settings. Maximum ingress-bandwidth takes the packet preamble and inter-frame gap (total of 20 bytes/packet) into account when counting a packet's bytes against the rate limiting meter whereas maximum egress bandwidth doesn’t. | Use policy based maximum bandwidth to limit ingress bandwidth.<br><br>->policy condition *irate* source port 1/1/1<br><br>->policy action *irate* maximum bandwidth 100M<br><br>->policy rule *irate* condition *irate* action *irate* |

| | | |
|---|---|---|
| 227159 | 9900 and 6560 maximum ingress depth is not working. | There is no known workaround at this time. |
| 227203 | Care must be taken when configuring maximum ingress-depth settings on a port, since setting this to too large a value can allow long high-speed bursts of packets to enter the switch. If these bursts are destined to a slower egress port, they may overflow the queuing capacity of that port and cause packet drops which will appear as a lower flow rate through that port.<br><br>Ingress rate limiting is based on an instantaneous policer, which simply allows packets through with their original inter-frame timing. Egress port rate limiting changes the scheduling rate for that port's egress queues and reshapes the traffic flow to have an inter-frame timing based on that rate. Because of this bursts of packets have to be queued on egress if the port speed it set to less than the burst rate of the ingress port and it's this queuing capacity that can be exceeded for long bursts of packets. | The switch will automatically calculate a default burst size based on the ingress rate that is set to try and guarantee that rate and limit bursts relative to that rate. It is recommended to not set the `maximum ingress-depth' setting unless traffic is fundamentally bursty and packet drops have been noticed due to the default setting.<br><br>When setting the maximum-depth take into account any speed differences between the ingress port and destination of that traffic to avoid egress queue drops due to packet bursts.<br><br>Additionally, it is highly recommended to also rate limit the sending link when setting an ingress rate limit. Since the ingress limit is a simple policer, exceeding the burst limit simply drops whatever packet is coming in without regard to priority. Egress rate shaping typically involves priority queuing, so will not only smooth out packet bursts, it will also be able to prioritize which packets to drop if the traffic rate is exceeded. |
| 227370 | With a QoS rule applied to rate-limit the ingress traffic and after ARP-timeout the switch sends an ARP request and end-user sends ARP-Reply. The switch is dropping that ARP-Reply and the traffic which is egressing out of the rate-limited port gets dropped. | There is no known workaround at this time. |
| 227421 | QoS stats rules take more time to display through WebView and some errors may be displayed on the console. | There is no known workaround at this time. |
| 228383 | qosCmm Config error message is seen on console after bootup on an OS6560. | There is no known workaround at this time. There is no functional impact. |
| 228452 | When an invalid policy validity period is entered the process exits and displays errors. | Re-enter a correct range for the validity period. |

**Authentication / UNP**

| PR | Description | Workaround |
|---|---|---|
| 227321 | UNP user traffic is getting dropped after NI reset. | There is no known workaround at this time. |
| 227684 | UNP rate limiting doesn't work on 6560 and 9900. | There is no known workaround at this time. |
| 227790 | UNP access port doesn't support following features.<br><br>1) Session-timeout,<br><br>2) Inactivity-logout<br><br>3) Accountint-Interim-interval | There is no known workaround at this time. |
| 228274 | When the UNP user is learned in built-in role due to LTP failure, the traffic from the UNP user should be dropped until the user moves out of LTP failure but it is not. | There is no known workaround at this time. |

**Layer 2**

| PR | Description | Workaround |
|---|---|---|
| 227836 | STP packets are not tunneled when service l2 profile 'unp-def-access-profile' is enabled on an UNP port and UNP Linkagg. | There is no known workaround at this time. |
| 225000 | ERP status does not change if service vlan is removed while ERP configurations are still intact. | Remove the ERP configuration before removing the VLAN association. |

**Virtual Chassis / Takeover**

| PR | Description | Workaround |
|---|---|---|
| 227989 | With VC-takover QoS stats are not getting updated until VC-takover CMM comes up. | Wait for the takeover to complete. |
| 227994 | With VC-takover on QoS config , errors are being displayed. | There is no known workaround at this time. |
| 228215 | 40G VFL configuration may cause system instability and malfunctioning of VC. | Recommended not to use 40G VFL configuration. Use 10G VFL. |
| 228350 | ISSU support for SPB-M is not available from 841.R01 to 841.R02. | There is no known workaround at this time. |

| | | |
|---|---|---|
| 228406 | On an OS9900 VC-2, after VC-takover the "copy running certified" command is getting an error for Secondary Synchronization. | Retry the command. |
| 228412 | svcCmm mDREQ & svcCmm mIPMS error messages are seen on console after VC-takeover & mac-flush. | There is no known workaround at this time. There is no functional impact. |
| 228715 | 9900 VC of 2 may cause system instability when configured with single CMM in each VC node. | Install redundant CMM on each chassis with 9900 VC of 2 configuration. |

## Hot Swap/Redundancy Feature Guidelines

### Hot Swap Feature Guidelines

Refer to the table below for hot swap/insertion compatibility. If the modules are not compatible a reboot of the chassis is required after inserting the new module.

- When connecting or disconnecting a power supply to or from a chassis, the power supply must first be disconnected from the power source.
- For the OS6900-X40 wait for first module to become operational before adding the second module.
- 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.

| Existing Expansion Slot | Hot-swap/Hot-insert compatibility |
|---|---|
| Empty | OS-XNI-U12, OS-XNI-U4 |
| OS-XNI-U4 | OS-XNI-U12, OS-XNI-U4 |
| OS-XNI-U12 | OS-XNI-U12, OS-XNI-U4 |
| OS-HNI-U6 | OS-HNI-U6 |
| OS-QNI-U3 | OS-QNI-U3 |
| OS-XNI-T8 | OS-XNI-T8 |
| OS-XNI-U12E | OS-XNI-U12E |

**OS6900 Hot Swap/Insertion Compatibility**

| Existing Slot | Hot-swap/Hot-insert compatibility |
|---|---|
| Empty | All modules can be inserted |
| OS99-CMM | OS99-CMM |
| OS9907-CFM | OS9907-CFM |
| OS99-GNI-48 | OS99-GNI-48 |
| OS99-GNI-P48 | OS99-GNI-P48 |
| OS99-XNI-48 | OS99-XNI-48 |
| OS99-XNI-U48 | OS99-XNI-U48 |
| OS99-GNI-U48 | OS99-GNI-U48 |

**OS9900 Hot Swap/Insertion Compatibility**

## 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 a replacement.
4. Insert replacement module of same type.
5. Follow any messages that may displayed.
6. Re-insert all transceivers into the new module.
7. Re-connect all cables to transceivers.
8. Hot swap one CFM at a time. Please ensure all fan trays are always inserted and operational. CFM hot swap should be completed with 120 seconds.

## 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 |
| European Union | +800 00200100 (Toll Free) or +1(650)385-2193 |
| Asia Pacific | +65 6240 8484 |

**Email:** ebg_global_supportcenter@al-enterprise.com

**Internet:** Customers with service agreements may open cases 24 hours a day via the support web page at: support.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 hardware configuration, module types 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.

## Third Party Licenses and Notices

Legal Notices applicable to any software distributed alone or in connection with the product to which this document pertains, are contained in files within the software itself located at: **/flash/foss**.

## Appendix A: Feature Matrix

The following is a feature matrix for AOS Release 8.4.1.R02.

***Note: Early availability features are available in AOS and can be configured. However, they have not gone through the complete AOS validation cycle and are therefore not officially supported.***

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| **Management Features** | | | | | | |
| USB Console Support | N | Y | N | N | N | |
| SNMP v1/v2/v3 | Y | Y | Y | Y | Y | |
| NTP | Y | Y | Y | Y | Y | |
| PING and TRACEROUTE as a Read-Only user | Y | Y | Y | Y | Y | |
| USB Disaster Recovery | Y | Y | Y | Y | Y | |
| Automatic Remote Configuration / Zero touch provisioning | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| IP Managed Services | Y | Y | Y | Y | Y | |
| SSH for read-only users | Y | Y | Y | Y | Y | |
| VRF | Y | Y | Y | N | Y | |
| VRF – DHCP Client | Y | Y | Y | N | Y | |
| Automatic/Intelligent Fabric | Y | Y | Y | Y | Y | |
| Automatic VC | Y | Y | Y | Y | Y | |
| Bluetooth for Console Access | N | Y | N | N | N | |
| EEE support | Y | Y | Y | Y | Y | |
| Embedded Python Scripting / Event Manager | Y | Y | Y | Y | Y | |
| ISSU | Y | Y | Y | Y | Y | |
| OpenFlow | Y | Y | N | N | N | |
| SAA | Y | Y | Y | Y | N | |
| SNMPv3 FIPS Certified Cryptographic Algorithms | N | N | N | N | N | |
| UDLD | Y | Y | Y | Y | N | |
| USB Flash | Y | Y | Y | Y | N | |
| Virtual Chassis (VC) | Y | Y | Y | Y | Y | |
| VC Split Protection (VCSP) | Y | Y | Y | Y | N | |
| Remote Chassis Detection (RCD) | Y | N | N | N | N | |
| Web Services & CLI Scripting | Y | Y | Y | Y | Y | |
| | | | | | | |
| Layer 3 Feature Support | | | | | | |
| ARP | Y | Y | Y | Y | Y | |
| OSPFv2 | Y | Y | Y | N | Y | |
| Static routing to an IP interface name | Y | Y | Y | Y | Y | |

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| ECMP | Y | Y | Y | Y | Y | |
| IGMP v1/v2/v3 | Y | Y | Y | Y | Y | |
| PIM-DM | Y | Y | Y | N | Y | |
| IPv4 Multicast Switching | Y | Y | Y | Y | Y | |
| Add tags to static-route command to enable easier redistribution | Y | Y | Y | Y | Y | |
| BGP with graceful restart | Y | Y | Y | N | Y | Not supported on a VC of 2 on the OS9900. |
| BGP route reflector for IPv6 | Y | Y | Y | N | N | |
| BGP ASPATH Filtering for IPv6 routes on IPv6 peering | Y | Y | Y | N | N | |
| BGP support of MD5 password for IPv6 | Y | Y | Y | N | N | |
| BGP 4-Octet ASN Support | Y | Y | Y | N | Y | Not supported on a VC of 2 on the OS9900. |
| GRE | Y | Y | Y | N | N | |
| IP-IP tunneling | Y | Y | Y | N | N | |
| IP routed port | Y | Y | Y | Y | Y | |
| IPv6 | Y | Y | Y | N | Y | JITC mode only on OS9900 |
| IPv6 DHCP relay and Neighbor discovery proxy | Y | Y | Y | N | N | |
| ISIS IPv4/IPv6 | Y | Y | Y | N | Y | JITC mode only on OS9900 |
| M-ISIS | Y | Y | Y | N | N | |
| OSPFv3 | Y | Y | Y | N | Y | JITC mode only on OS9900 |
| RIP v1/v2 | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| RIPng | Y | Y | Y | N | N | |
| DHCP Server (v4, v6 with integrated support of QIP remote management) | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| VRRP v2 | Y | Y | Y | Y | Y | |
| VRRP v3 | Y | Y | Y | N | N | |
| ARP - Proxy | Y | Y | Y | Y | Y | |
| ARP - Distributed | Y | N | N | N | N | |
| BFD | Y | Y | Y | N | Y | |
| DHCP Snooping | Y | Y | Y | Y | Y | |
| DHCP Snooping IP source filtering – VLAN/port-based | Y | Y | Y | Y | N | |
| DHCPv6 Relay | Y | Y | Y | Y | Y | |
| IP Multinetting | Y | Y | Y | Y | Y | |

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| IPSec | Y | Y | Y | N | Y | JITC mode only on OS9900 |
| Server Load Balancing (SLB) | Y | Y | Y | N | N | |
| | | | | | | |
| Multicast Features | | | | | | |
| IGMP v1/v2/v3 | Y | Y | Y | Y | Y | |
| IPv4 Multicast Switching | Y | Y | Y | Y | Y | |
| PIM-DM | Y | Y | Y | N | Y | Not supported on a VC of 2 on the OS9900. |
| DVMRP | Y | Y | Y | N | N | |
| IPv6 Multicast Switching (MLD v1/v2) | Y | Y | Y | N | N | |
| IPv6 Scoped Multicast Addresses | Y | Y | Y | N | N | |
| PIM-SM | Y | Y | Y | N | Y | |
| PIM-SSM | Y | Y | Y | N | Y | |
| PIM-SSM Static Map | Y | Y | Y | N | N | |
| PIM-BiDir | Y | Y | Y | N | Y | Not supported on a VC of 2 on the OS9900. |
| | | | | | | |
| Monitoring/Troubleshooting Features | | | | | | |
| Extended ping and traceroute | Y | Y | Y | Y | Y | |
| Port mirroring | Y | Y | Y | Y | Y | |
| Port monitoring | Y | Y | Y | Y | Y | |
| Switch logging / Syslog | Y | Y | Y | Y | Y | |
| RMON | Y | Y | Y | Y | Y | |
| SFlow | Y | Y | Y | N | N | |
| Policy based mirroring | Y | Y | Y | Y | N | |
| Port mirroring - remote | Y | Y | Y | Y | N | |
| TDR | N | Y | N | N | N | |
| | | | | | | |
| Layer 2 Feature Support | | | | | | |
| 802.1q | Y | Y | Y | Y | Y | |
| Spanning Tree (802.1ad, 802.1w, MSTP, PVST+, Root Guard) | Y | Y | Y | Y | N | |
| LLDP (802.1ab) | Y | Y | Y | Y | Y | |
| Link Aggregation (static and LACP) | Y | Y | Y | Y | Y | |
| STP Loop Guard | Y | Y | Y | Y | Y | |
| DHL | N | Y | Y | Y | N | |
| ERP v1/v2 | Y | Y | Y | N | N | |

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| HAVLAN | Y | Y | Y | N | N | |
| Loopback detection – Edge (Bridge) | N | Y | Y | N | Y | |
| Loopback detection – SAP (Access) | Y | Y | Y | N | N | |
| MVRP | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| Private VLANs | Y | Y | Y | N | N | |
| Source Learning – Distributed Mode | N | N | N | N | N | |
| SIP Snooping | N | Y | N | N | N | |
| | | | | | | |
| QoS Feature Support | | | | | | |
| QSP Profiles | Y | Y | Y | Y | Y | |
| Per port rate limiting | Y | Y | Y | Y | Y | |
| 802.1p / DSCP priority mapping | Y | Y | Y | Y | Y | |
| Auto-Qos prioritization of NMS/IP Phone Traffic | Y | Y | Y | Y | Y | |
| ACL – IPv4 | Y | Y | Y | Y | Y | |
| ACL – IPv6 | Y | Y | Y | Y | N | |
| MAC Groups | Y | Y | Y | Y | Y | |
| Network Groups | Y | Y | Y | Y | Y | |
| Port Groups | Y | Y | Y | Y | Y | |
| Service Groups | Y | Y | Y | Y | Y | |
| Map Groups | Y | Y | Y | Y | Y | |
| Switch Groups | Y | Y | Y | Y | Y | |
| Policy Lists | Y | Y | Y | Y | Y | |
| Policy based routing | Y | Y | Y | Y | Y | |
| Ingress/Egress bandwidth limit | Y | Y | Y | Y | Y | |
| Tri-color marking | Y | Y | Y | Y | N | |
| QSP Profiles 2/3/4 | Y | Y | Y | Y | N | QSP 1 and 5 supported on the OS9900. |
| | | | | | | |
| Metro Ethernet Features | | | | | | |
| Ethernet Services | Y | Y | Y | N | N | |
| Ethernet OAM (ITU Y1731 and 802.1ag) | Y | Y | Y | N | N | |
| | | | | | | |
| Security Features | | | | | | |
| Access Guardian – Bridge | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| Access Guardian - Access | N | Y | Y | N | N | |
| Interface Violation Recovery | Y | Y | Y | Y | Y | |
| Learned Port Security (LPS) | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| LLDP Rogue Detection | Y | Y | Y | Y | Y | Not supported on a VC of 2 on the OS9900. |
| TACACS+ Client | Y | Y | Y | Y | Y | |
| TACACS+ command based authorization | Y | Y | Y | N | Y | |
| Accounting | Y | Y | Y | Y | Y | |
| Application Monitoring and Enforcement (Appmon) | N | Y | N | N | N | |
| ARP Poisoning Protection | Y | Y | Y | Y | Y | |
| Application Fingerprinting | Y | N | N | N | N | |
| COA Extension support for RADIUS (BYOD) | N | Y | Y | Y | Y | |
| mDNS Snooping/Relay (BYOD) | N | Y | Y | Y | Y | |
| UPNP/DLNA Relay (BYOD) | N | Y | Y | Y | Y | |
| Switch Port location information pass-through in RADIUS requests (BYOD) | N | Y | Y | Y | Y | |
| Captive Portal | N | Y | Y | Y | Y | |
| Quarantine Manager | N | Y | Y | Y | Y | |
| Radius test tool | Y | Y | Y | Y | Y | |
| Storm Control | Y | Y | Y | Y | N | |
| | | | | | | |
| PoE Features | | | | | | |
| 802.1af and 802.3at | N | Y | Y | Y | Y | |
| Auto Negotiation of PoE Class-power upper limit | N | Y | Y | Y | Y | |
| Display of detected power class | N | Y | Y | Y | Y | |
| LLDP/802.3at power management TLV | N | Y | Y | Y | Y | |
| HPOE support | N | Y (60W) | Y (75W) | Y (95W) | Y (75W) | |
| POE Time Of Day Support | N | Y | Y | Y | Y | |
| | | | | | | |
| Data Center Features | | | | | | |
| CEE DCBX Version 1.01 | Y | N | N | N | N | |
| Data Center Bridging (DCBX/ETS/PFC) | Y | N | N | N | N | |

| Feature | OS6900 | OS6860(E) | OS6865 | OS6560 | OS9900 | Notes |
|---|---|---|---|---|---|---|
| EVB | Y | N | N | N | N | |
| FCoE / FC Gateway | Y | N | N | N | N | |
| FIP Snooping | Y | N | N | N | N | |
| IPv4 over SPB | Y | Y | Y | N | N | |
| RFP on SPB UNI port | Y | N | N | N | N | |
| SPB | Y | Y | Y | N | N | |
| VXLAN | Q32/X72 | N | N | N | N | |
| VM/VXLAN Snooping | Y | N | N | N | N | |
| | | | | | | |

## Appendix B: General Upgrade Requirements and Best Practices

This section is to assist with upgrading an OmniSwitch. The goal is to provide a clear understanding of the steps required and to answer any questions about the upgrade process prior to upgrading. Depending upon the AOS version, model, and configuration of the OmniSwitch various upgrade procedures are supported.

**Standard Upgrade** - The standard upgrade of a standalone chassis or virtual chassis (VC) is nearly identical. All that's required is to upload the new image files to the *Running* directory and reload the switch. In the case of a VC, prior to rebooting the Master will copy the new image files to the Slave(s) and once the VC is back up the entire VC will be synchronized and running with the upgraded code.

**ISSU** - The In Service Software Upgrade (ISSU) is used to upgrade the software on a VC or modular chassis with minimal network disruption. Each element of the VC is upgraded individually allowing hosts and switches which are dual-homed to the VC to maintain connectivity to the network. The actual downtime experienced by a host on the network should be minimal but can vary depending upon the overall network design and VC configuration. Having a redundant configuration is suggested and will help to minimize recovery times resulting in sub-second convergence times.

**Virtual Chassis** - The VC will first verify that it is in a state that will allow a successful ISSU upgrade. It will then copy the image and configuration files of the ISSU specified directory to all of the Slave chassis and reload each Slave chassis from the ISSU directory in order from lowest to highest chassis-id. For example, assuming chassid-id 1 is the Master, the Slave with chassis-id 2 will reload with the new image files. When Slave chassis-id 2 has rebooted and rejoined the VC, the Slave with chassis -id 3 will reboot and rejoin the VC. Once the Slaves are complete they are now using the new image files. The Master chassis is now rebooted which causes the Slave chassis to become the new Master chassis. When the original Master chassis reloads it comes back as a Slave chassis. To restore the role of Master to the original Master chassis the current Master can be rebooted and the original Master will takeover, re-assuming the Master role.

**Modular Chassis** - The chassis will first verify that it is in a state that will allow a successful ISSU upgrade. It will then copy the image and configuration files of the ISSU specified directory to the secondary CMM and reload the secondary CMM which becomes the new primary CMM. The old primary CMM becomes the secondary CMM and reloads using the upgraded code. As a result of this process both CMMs are now running with the upgraded code and the primary and secondary CMMs will have changed roles (i.e., primary will act as secondary and the secondary as primary). The individual NIs can be reset either manually or automatically (based on the NI reset timer).

## Supported Upgrade Paths and Procedures

The following releases support upgrading using ISSU. All other releases support a Standard upgrade only.

| Platform | AOS Releases Supporting ISSU to 8.4.1.R02 (GA) |
|---|---|
| OS6900 | 8.3.1.314.R01 (GA)<br>8.3.1.160.R02 (GA)<br>8.3.1.180.R02 (MR)<br>8.4.1.170.R01 (GA) |
| OS6860(E) | 8.4.1.170.R01 (GA) |
| OS6865 | 8.3.1.314.R01 (GA)<br>8.3.1.160.R02 (GA)<br>8.3.1.180.R02 (MR)<br>8.4.1.170.R01 (GA) |
| OS6560 | Not Supported |
| OS9900 | Not Supported<br>**Note:** Due to architectural changes required for VC of 2 support. ISSU is not supported from previous releases. |
| **Note:** All GA hardware has been released with the proper Uboot and FPGA version. There are currently no Uboot or FPGA upgrades required.<br><br>**Note**: ISSU is not support on any platforms that are running IPMS over SPB. Use a standard upgrade. | |

**ISSU Supported Releases**

## Prerequisites

These upgrade instructions require 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 access may be affected by following 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 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 and FPGA. If they meet the minimum requirements, (i.e. they were already upgraded during a previous AOS upgrade) then only an upgrade of the AOS images is required.
- Depending on whether a standalone chassis or VC is being upgraded, upgrading can take from 5 to 20 minutes. Additional time will be needed for the network to re-converge.

- The examples below use various models and directories to demonstrate the upgrade procedure. However any user-defined directory can be used for the upgrade.

- If possible, have EMP or serial console access to all chassis during the upgrade. This will allow you to access and monitor the VC during the ISSU process and before the virtual chassis has been re-established.

- Knowledge of various aspects of AOS directory structure, operation and CLI commands can be found in the Alcatel-Lucent OmniSwitch User Guides. Recommended reading includes:
    - Release Notes - for the version of software you're planning to upgrade to.
    - The AOS Switch Management Guide
        - Chapter – Getting Started
        - Chapter - Logging Into the Switch
        - Chapter - Managing System Files
        - Chapter - Managing CMM Directory Content
        - Chapter - Using the CLI
        - Chapter - Working With Configuration Files
        - Chapter - Configuring Virtual Chassis

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

### Switch Maintenance

It's recommended to perform switch maintenance prior to performing any upgrade. This can help with preparing for the upgrade and removing unnecessary files. The following steps can be performed at any time prior to a software upgrade. These procedures can be done using Telnet and FTP, however using SSH and SFTP/SCP are recommended as a security best-practice since Telnet and FTP are not secure.

1. Use the command '**show system**' to verify current date, time, AOS and model of the switch.

```
6900-> show system
System:
  Description:  Alcatel-Lucent OS6900-X20 7.3.2.568.R01 Service Release, September 05, 2014.,
  Object ID:    1.3.6.1.4.1.6486.801.1.1.2.1.10.1.1,
  Up Time:      0 days 0 hours 1 minutes and 44 seconds,
  Contact:      Alcatel-Lucent, http://alcatel-lucent.com/wps/portal/enterprise,
  Name:         6900,
  Location:     Unknown,
  Services:     78,
  Date & Time:  FRI OCT 31 2014 06:55:43 (UTC)
Flash Space:
    Primary CMM:
      Available (bytes):  1111470080,
      Comments         :  None
```

2. Remove any old tech_support.log files, tech_support_eng.tar files:

```
6900-> rm *.log
6900-> rm *.tar
```

3. Verify that the **/flash/pmd** and **/flash/pmd/work** directories are empty. If they have files in them check the date on the files. If they are recently created files (<10 days), contact Alcatel-Lucent Service & Support. If not, they can be deleted.

4. Use the '**show running-directory**' command to determine what directory the switch is running from and that the configuration is certified and synchronized:

```
6900-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : vc_dir,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Running Configuration    : SYNCHRONIZED
```

If the configuration is not certified and synchronized, issue the command '**write memory flash-synchro**':

```
6900-> write memory flash-synchro
```

6. If you do not already have established baselines to determine the health of the switch you are upgrading, now would be a good time to collect them. Using the show tech-support series of commands is an excellent way to collect data on the state of the switch. The show tech support commands automatically create log files of useful show commands in the **/flash** directory. You can create the tech-support log files with the following commands:

```
6900-> show tech-support
6900-> show tech-support layer2
6900-> show tech-support layer3
```

Additionally, the '**show tech-support eng complete**' command will create a TAR file with multiple tech-support log files as well as the SWLOG files from the switches.

```
6900-> show tech-support eng complete
```

It is a good idea to offload these files and review them to determine what additional data you might want to collect to establish meaningful baselines for a successful upgrade.

- If upgrading a standalone chassis or VC using a standard upgrade procedure please refer to Appendix C for specific steps to follow.
- If upgrading a VC using ISSU please refer to Appendix D for specific steps to follow.

## Appendix C: Standard Upgrade - OmniSwitch Standalone or Virtual Chassis

These instructions document how to upgrade a standalone or virtual chassis using the standard upgrade procedure. Upgrading using the standard upgrade procedure consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go to the Service and Support website and download and unzip the upgrade files for the appropriate model and release. The archives contain the following:

- OS6900 - Tos.img
- OS6860 – Uos.img
- OS6865 – Uos.img
- OS6560 – Nos.img
- OS9900 – Mos.img, Mhost.img, Meni.img
- imgsha256sum (not required) –This file is only required when running in Common Criteria mode. Please refer to the Common Criteria Operational Guidance Document for additional information. (**Note**: This document will be available at a future date after completion of Common Criteria certification).

2. FTP the Upgrade Files to the Switch

FTP the image files to the *Running* directory of the switch you are upgrading. The image files and directory will differ depending on your switch and configuration.

3. Upgrade the image file

Follow the steps below to upgrade the image files by reloading the switch from the *Running* directory.

```
OS6900-> reload from working no rollback-timeout
Confirm Activate (Y/N) : y
This operation will verify and copy images before reloading.
It may take several minutes to complete....
```

If upgrading a VC the new image file will be copied to all the Slave chassis and the entire VC will reboot. After approximately 5-20 minutes the VC will become operational.

4. Verify the Software Upgrade

Log in to the switch to confirm it is running on the new software. This can be determined from the login banner or the **show microcode** command**.**

```
OS6900-> show microcode
   /flash/working
   Package           Release                 Size     Description
-----------------+-----------------------+--------+-----------------------------------
Tos.img           8.4.1.229.R02             210697424 Alcatel-Lucent OS
```

```
-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : WORKING,
  Certify/Restore Status   : CERTIFY NEEDED
SYNCHRONIZATION STATUS
  Running Configuration    : SYNCHRONIZED
```

**Note**: If there are any issues after upgrading the switch can be rolled back to the previous certified version by issuing the **reload from certified no rollback-timeout** command.

5. Certify the Software Upgrade

After verifying the software and that the network is stable, use the following commands to certify the new software by copying the *Running* directory to the Certified directory.

```
OS6900-> copy running certified
Please wait…….

-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : WORKING,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Running Configuration    : SYNCHRONIZED
```

## Appendix D: ISSU – OmniSwitch Chassis or Virtual Chassis

These instructions document how to upgrade a modular chassis or virtual chassis using ISSU. Upgrading using ISSU consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go to the Service and Support Website and download and unzip the ISSU upgrade files for the appropriate platform and release. The archive contains the following:

- OS6900 **-** Tos.img
- OS6860 – Uos.img
- OS6865 – Uos.img
- OS6560 – Nos.img
- OS9900 – Mos.img, Mhost.img, Meni.img
- ISSU Version File – issu_version
- imgsha256sum (not required) –This file is only required when running in Common Criteria mode. Please refer to the Common Criteria Operational Guidance Document for additional information. (**Note**: This document will be available at a future date after completion of Common Criteria certification).

**Note:** The following examples use **issu_dir** as an example ISSU directory name. However, any directory name may be used. Additionally, if an ISSU upgrade was previously performed using a directory named **issu_dir**, it may now be the *Running Configuration*, in which case a different ISSU directory name should be used.

2. Create the new directory on the Master for the ISSU upgrade:

```
OS6900-> mkdir /flash/issu_dir
```

3. Clean up existing ISSU directories

It is important to connect to the Slave chassis and verify that there is no existing directory with the path **/flash/issu_dir** on the Slave chassis. ISSU relies upon the switch to handle all of the file copying and directory creation on the Slave chassis. For this reason, having a pre-existing directory with the same name on the Slave chassis can have an adverse affect on the process. To verify that the Slave chassis does not have an existing directory of the same name as the ISSU directory on your Master chassis, use the internal VF-link IP address to connect to the Slave. In a multi-chassis VC, the internal IP addresses on the Virtual Fabric Link (VFL) always use the same IP addresses: 127.10.1.65 for Chassis 1,127.10.2.65 for Chassis 2, etc. These addresses can be found by issuing the debug command '**debug show virtual-chassis connection**' as shown below:

```
OS6900-> debug show virtual-chassis connection
                                    Address             Address
 Chas  MAC-Address          Local IP             Remote IP          Status
-----+------------------+--------------------+------------------+-------------
 1       e8:e7:32:b9:19:0b  127.10.2.65       127.10.1.65       Connected
```

4. SSH to the Slave chassis via the internal virtual-chassis IP address using the password 'switch':

```
OS6900-> ssh 127.10.2.65

Password:switch
```

5. Use the **ls** command to look for the directory name being used for the ISSU upgrade. In this example, we're using **/flash/issu_dir** so if that directory exists on the Slave chassis it should be deleted as shown below. Repeat this step for all Slave chassis:

```
6900-> rm -r /flash/issu_dir
```

6. Log out of the Slave chassis:

```
6900-> exit
logout
Connection to 127.10.2.65 closed.
```

7. On the Master chassis copy the current *Running* configuration files to the ISSU directory:

```
OS6900-> cp /flash/working/*.cfg /flash/issu_dir
```

8. FTP the new image files to the ISSU directory. Once complete verify that the ISSU directory contains only the required files for the upgrade:

```
6900-> ls /flash/issu_dir
Tos.img       issu_version  vcboot.cfg    vcsetup.cfg
```

9. Upgrade the image files using ISSU:

```
OS6900-> issu from issu_dir
Are you sure you want an In Service System Upgrade? (Y/N) : y
```

During ISSU '**show issu status**' gives the respective status (pending, complete, etc)

```
OS6900-> show issu status
Issu pending
```

This indicates that the ISSU is completed

```
OS6900-> show issu status
Issu not active
```

Allow the upgrade to complete. DO NOT modify the configuration files during the software upgrade. It normally takes between 5 and 20 minutes to complete the ISSU upgrade. Wait for the System ready or [L8] state which gets displayed in the ssh/telnet/console session before performing any write-memory or configuration changes.

```
6900-> debug show virtual-chassis topology
Local Chassis: 1
Oper                                    Config   Oper                      System
Chas  Role         Status              Chas ID  Pri   Group  MAC-Address        Ready
-----+------------+-------------------+--------+-----+------+------------------+-------
1     Master        Running             1        100   19     e8:e7:32:b9:19:0b  Yes
2     Slave         Running             2        99    19     e8:e7:32:b9:19:43  Yes
```

10. Verify the Software Upgrade

Log in to the switch to confirm it is running on the new software. This can be determined from the login banner or the **show microcode** command.

```
OS6900-> show microcode
   /flash/working
   Package           Release                 Size     Description
-----------------+------------------------+--------+-----------------------------------
Tos.img           8.4.1.229.R02             210697424 Alcatel-Lucent OS
```

```
OS6900-> copy running certified
Please wait…….

-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : issu_dir,
  Certify/Restore Status   : CERTIFY NEEDED
SYNCHRONIZATION STATUS
  Flash Between CMMs       : SYNCHRONIZED
  Running Configuration    : SYNCHRONIZED
```

11. Certify the Software Upgrade

After verifying the software and that the network is stable, use the following commands to certify the new software by copying the *Running* directory to the Certified directory:

```
OS6900-> copy running certified
Please wait…….

-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : issu_dir,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Flash Between CMMs       : SYNCHRONIZED
  Running Configuration    : SYNCHRONIZED
```

## Appendix E: Fixed Problem Reports

The following problem reports were closed or are in verification in AOS Release 8.4.1.R02.

| PR | Summary |
|---|---|
| 215401 | Both Master and Master Split Topology holds the VC EMP Address |
| 218556 | Different OID's for AlcatelIND1Base.mib |
| 222557 | OS6860 Captive portal fails after 3 days. |
| 224249 | 'Show interface alias' command output display issue in a 3 unit Virtual chassis of OS6860E |
| 224483 | Wrong LACP Version errors in OS 6900 X20 |
| 224544 | Explanation for the log message udldNi Ni error(2) pktdrv_xmit_writebuffer:Could not send sys pdu -1 |
| 224560 | Fabric LED Of slave cmm stay in orange color and resloved after the reboot. |
| 225119 | OS6860: Clarification required on interface speed configuration for 10gig ports. |
| 225426 | OS6860 chassis 2 in a VC-7 reboot on Mar 27 17:29:45 and had a Kernel panic crash at Mar 27 17:29:46 |
| 225452 | mac-ping reports losses when intermediate SPB node has to transmit data over VFL |
| 225519 | "force-l3-learning" command is not saved in the vcboot.cfg on OS6860 |
| 225610 | 2XOS6860E(8.2.1.335.R01): No interfaces status shown for slave chassis |
| 225805 | OS6900 switches forwarding the DHCP Discover packets with the exit interface IP address instead of S |
| 225874 | The output of the "Show lldp remote-system med inv" command is mingled and causing confusion. |
| 225985 | OS6900-X72 : interface is going down if auto negotiation set as disabled. |
| 226105 | OS6860E: Network Time Protocol (NTP) Mode 6 queries are answered by AOS - vulnerability detected. |
| 226232 | [TYPE1]Vxlan multicast tunnel not masking with 24th bit of dst mac address |
| 226315 | OS6860: Need to have records on the swlogs for STP topology changes "Topology changed on VLAN/STP id |
| 226387 | OS9907 : 802.1x authentication failed. |
| 226415 | IP helper is not working in VRF |
| 226540 | OS6900: DoS type unicast dest-ip/multicast-mac in switch logs. |
| 226627 | OS6900- Display issue on "show interfaces" output |
| 226692 | 2XOS6860E VC: Issue on SNMP WALK with the OID .1.3.6.1.2.1.2.2.1.2 |
| 226741 | OSPF route convergence issue after configuring cost on OSPF interface |
| 226745 | OS6900-X40 Interface up - up without SFP |
| 226768 | OS9900: ISSU upgrade status shown incorrectly while performing the ISSU upgrade |
| 226917 | Customer has a VC of 6860, the master encounter a crash, log error "COREDUMPER alarm(1) Dumping cor |
| 227025 | DHL vlan  don't send out bogus mac frames in DHL raw mode on some vlan |
| 227117 | When MASTER EMP is configured, the loopback address becomes unreachable. A reboot is necessary for b |
| 227150 | 2xVC OS6900 : Discrepancies on the DEMO LICENSE expiry dates - increasing validity period after rebo |
| 227161 | OS6860-24 (841.170.R01): SNMP walk returns wrong value for the object ifNumber |
| 227198 | OS6900 VC - Frequent stpNi _TCHt logs |

| | |
|---|---|
| 227216 | AOS fails to send the response to CoA message sent by the radius server |
| 227235 | In a VC of 4, the master stays UP but all the slave crashed and rebooted. |
| 227344 | OS6865- Switch do not send power supply down trap to OV. |
| 227482 | OS6860 - Show UNP port range CLI is not working |
| 227523 | Displaying "ACCEPTED" message in swlogs even though logged in with incorrect SSH credentials |
| 227566 | OS6900 udprelaycmmd pmd file generated |
| 227571 | OS6860VC : OK1 LED is blinking after the AOS upgrade. |
| 227578 | OS6860 sending Remote Ip-address as 0.0.0.0 instead of User Ip-address in TACPLUS Auth Query packet. |
| 227586 | issue with configuring SNMP user through CLI |
| 227588 | OS9900 chassis report that there is not enough power supply to start the NI while we have enough pow |
| 227589 | OS9900 does not calculate the amount of power left in the output of show chassis command. |
| 227795 | OS6860 OIDs missing for NTP. |
| 227873 | Dynamic SPB path creation bug using UNP |
| 227885 | OS6860E-P24: policy rule for specific VRF is created; however, while executing the “show configurati |
| 227935 | Command needed to adjust the TACACS+ server timeout in 8x. |
| 227942 | OS6860E: MIB “Bridge-MIB dot1dBaseNumPorts” returns wrong value |
| 227949 | Switch port showing 'inactive' |
| 227980 | OS6900 internal dhcp server not working and reboot of switch fixes the issue |
| 228011 | OS9900: Error “error hal_mrSfpUtil_readSFPPhy” seen while connecting a SFP |
| 228013 | 6xOS6900: LLDP "ldpCmm Mgr warning(4) cmmLldpMipValidIfIndexRange" messages seen on console/ swlogs |
| 228156 | Issue of communication / synchronisation between BFD and OSPF. |
| 228167 | OS6860 interface speed change issue. |
| 228341 | OS6900: "lldpCmm Mgr warning” messages seen on the console/ swlogs |