Release Notes - Rev. D

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

Release 8.4.1.R03

These release notes accompany release 8.4.1.R03. 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

Contents 2

Related Documentation 3

System Requirements 4

[IMPORTANT] *MUST READ*: AOS Release 8.4.1.R03 Prerequisites and Deployment Information 6

Licensed Features 8

CodeGuardian 9

New / Updated Hardware Support 10

New Software Features and Enhancements 11

Open Problem Reports and Feature Exceptions 15

Hot Swap/Redundancy Feature Guidelines 18

Technical Support 20

Appendix A: Feature Matrix 21

Appendix B: General Upgrade Requirements and Best Practices 26

Appendix C: Standard Upgrade - OmniSwitch Standalone or Virtual Chassis 30

Appendix D: ISSU – OmniSwitch Chassis or Virtual Chassis 32

Appendix E: Fixed Problem Reports 35

Appendix F: PoE Firmware Upgrade Instructions 37

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.

- 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.R03 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 8.4.1.141.R03(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 8.4.1.141.R03(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 8.4.1.141.R03(GA)

Hardware	Minimum UBoot	Minimum FPGA
CMM All Expansion Modules	7.3.4.277.R01 N/A	0.1.8 N/A

OmniSwitch 6900-X72 – AOS Release 8.4.1.141.R03(GA)

Hardware	Minimum Uboot	Minimum FPGA
CMM All Expansion Modules	7.3.4.31.R02 N/A	0.1.10 N/A

OmniSwitch 6860(E) – AOS Release 8.4.1.141.R03(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)
OS6860E-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 8.4.1.141.R03(GA)

Hardware	Minimum Uboot	Minimum FPGA*
OS6865-P16X	8.3.1.125.R01	0.20 (0x14) (minimum) 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 8.4.1.141.R03(GA)

Hardware	Minimum Uboot	Minimum FPGA
OS6560-P24Z24	8.4.1.23.R02	0.6 (0x6)
OS6560-P24Z8	8.4.1.23.R02	-
OS6560-P48Z16	8.4.1.23.R02	0.6

OmniSwitch 9900 – AOS Release 8.4.1.141.R03(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
OS99-CNI-U8	8.4.1.20.R03	1.7	N/A
OS99-XNI-P48Z16	8.4.1.20.R03	1.4	0.6

[IMPORTANT] *MUST READ*: AOS Release 8.4.1.R03 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.R03 please refer to Appendix B for important best practices, prerequisites, and step-by-step instructions.

Additional Information

- All switches that ship from the factory with AOS Release 8.4.1.R03 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.R02 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.
- Beginning in AOS 8.4.1.R03 the **vrf *vrf-name*** command will no longer create a new VRF. The **vrf *vrf-name*** command only changes the CLI context to the VRF specified by ***vrf-name*** if the vrf-name exists.

The command **vrf create *vrf-name*** is used for creating a new VRF. When upgrading to 8.4.1.R03 all instances of the old syntax will be automatically converted to the new syntax.

- The "redirect" parameter used to configure BYOD access for UNP profiles has been deprecated. BYOD redirection for profile devices is now automatically available based on the status of Captive Portal authentication for the UNP profile.
    - When Captive Portal authentication is disabled for the profile, BYOD redirection is automatically made available to devices assigned to the profile. If the initial device authentication process returns the "Alcatel-Redirect-URL" attribute, then the device is automatically redirected to the BYOD server. No specific UNP profile configuration is required.
    - When Captive Portal authentication is enabled, internal Captive Portal is enforced and BYOD redirection is not available.
    - UNP profile redirect configuration is still allowed at boot up to support upgrade from previous releases.

OS9900 Virtual Chassis Guidelines

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)

OS99-GNI-P48 - 75 Watts of Power Over Ethernet (Possible PoE Firmware Upgrade Required)

Some OS99-GNI-P48 modules were shipped with a PoE firmware version that supports a maximum of 60 watts of PoE on the HPoE ports instead of the the 75 watts of PoE documented. If 75 watts of PoE is required refer to the table below and follow the instructions in Appendix F if a PoE firmware upgrade is required.

OS99-GNI-P48 Board Revision	PoE Firmware Version Shipped with Module
Rev C09 or lower	00.0263.01 (Supports maximum of 60W of PoE on HPoE ports)
Rev C10 or higher	00.0265.04 (Supports maximum of 75W of PoE on HPoE ports)

Use the '**show module**' command to display the current module revision. If the current revision is lower than C10 the PoE firmware needs to be upgraded to support 75 watts.

```
-> show module 4
                                HW
Chassis/Slot     Part-Number      Serial #    Rev    Mfg Date    Model Name
---------------+---------------+-------------+----+-------------+--------------
 1/SLOT-4       903726-90       U1920084      C05   Jun 2 2016   OS99-GNI-P48
```

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?
	OmniSwitch 6900
Data Center Features	
DCB (PFC,ETS,DCBx)	Yes
EVB	Yes
FIP Snooping	Yes
FCoE VXLAN	Yes
Note: All other platforms do not support Data Center features.	

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.

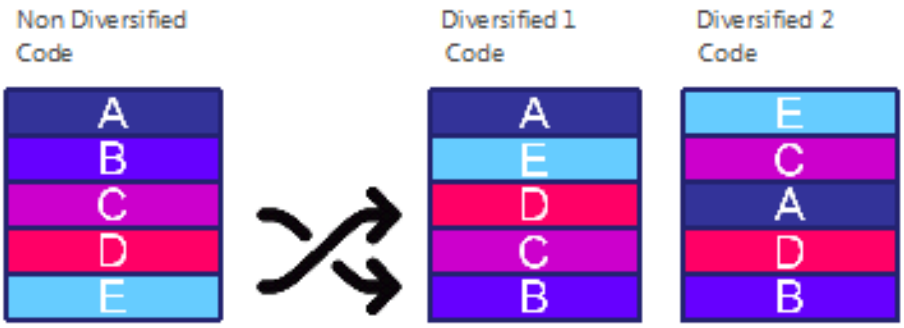

CodeGuardian AOS Releases

Standard AOS Releases	AOS CodeGuardian Release	LGS AOS CodeGuardian Release
AOS 8.4.1.R03	AOS 8.4.1.RX3	AOS 8.4.1.LX3

- X=Diversified image 1-3
- ALE will have 3 different diversified images per AOS release (R11 through R31)
- Our partner LGS will have 3 different diversified images per AOS release (L11 through L31)

Please contact customer support for additional information.

New / Updated Hardware Support

The following new hardware is being introduced in this release.

OS99-XNI-P48Z16

The OS99-XNI-P48Z16 is a multi-gig network interface card for the OS9900 offering:

- Thirty-two (32) – 1G/10G Base-T 802.3at PoE ports (30W PoE)
- Sixteen (16) – 1G/2.5G/10G BaseT 802.3bt PoE ports (Ports 1-8 support 75W HPoE)

OS99-CNI-U8

The OS99-CNI-U8 is a 100G network interface card for the OS9900 offering:

- Eight (8) – 100G unpopulated wire rate QSFP28 ports.

OS6560-P48Z16

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

Transceivers

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

New Transceivers	Platform Support
40G Transceivers	
QSFP-40G-CLR (future availability)	OS6900-Q32, OS6900-QNI-U3, OS99-CMM
QSFP-40G-LM4 (future availability)	OS6900-Q32, OS6900-QNI-U3, OS99-CMM
100G Transceivers	
QSFP-100G-C1M QSFP-100G-C3M QSFP-100G-C5M	OS99-CNI-U8
QSFP-100G-A20M	OS99-CNI-U8
QSFP-100G-CLR4	OS99-CNI-U8
QSFP-100G-CWDM4	OS99-CNI-U8
QSFP-100G-LR4	OS99-CNI-U8
QSFP-100G-SR4	OS99-CNI-U8

- The OS99-CNI-U8 does not support auto-negotiation on 100G ports.
- 100G fiber transceivers are supported between an OmniSwitch and other vendor switches, but the remote switch must disable auto-negotiation.

New Software Features and Enhancements

The following software features are being introduced with the 8.4.1.R03 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.R03 New Feature/Enhancements Summary

Feature	Platform
Remote Port Mirroring to Linkagg	OS6900/6860(E)/6865
VC of 2 Enhancements on 9900	OS9900
L3 VPN Interface / SPB-M IP Inline Routing	OS9900
SPB-M 9900 support	OS9900
SPB Auto-fabric 9900 Support	OS9900
VRF "create" Command	OS6900/6860(E)/6865/9900
IPv6 Enhancements for the OS9900: • IPv6 VRF • RIPng • OSPFv3 • BGP4 IPv6 Extension • IPv6 Unicast Basics • IPv6 MC Routing • IPv6 DHCP Relay • VRRPv3 • IPv6 MC Switching	OS9900
BFD Support	OS9900
OmniVista Cirrus (OV Cloud)	OS6900/6860(E)/6865/6560
MACSec	OS6860(E)
1588 Transparent Clock	OS6860(E)/6865
ZeroConf mDNS and SSDP Enhancements	All
Four/two-pair Power Over Ethernet	OS6860(E)/6865/6560/9900

Remote Port Mirroring to Link Aggregate
Allows remote port mirroring to a link aggregate.

VC of Two Enhancements on OS9900
Refer to feature matrix for a list of OS9900 features that are not supported on an OS9900 VC of 2.

L3 VPN Interface - SPB-M IP Inline Routing

The OmniSwitch IP over SPBM solution supports two methods for routing L3 IP traffic over an L2 SPBM network: VPN-Lite and L3 VPN. Both of these methods require an L3 VPN interface to serve as an IP gateway for accessing remote networks.

There are two options for defining an L3 VPN interface:

- Configuring a service-based IP interface (OmniSwitch 9900 only). When this option is used, an IP interface is created and bound to an SPB service ID. The service ID is associated with an I-SID.
- Configuring an external loopback port configuration (Not supported on the OS9900). When this option is used, a physical cable connects a regular port to a service access port. The regular port is tagged with an IP interface VLAN; the access port is associated with an SPB Service Access Point (SAP). The VLAN-based IP interface serves as the L3 VPN interface.

The external loopback configuration option defines an L3 VPN loopback interface, and the service-based IP interface option defines an L3 VPN service-based interface.

L3 VPN Service-Based Interface

Configuring an L3 VPN loopback interface and/or an L3 VPN service-based interface is supported on an OmniSwitch 9900. Using the service-based interface option simplifies the configuration in that a physical cable and two switch ports are not required to form an external loopback. Instead, an IP interface is created within a VRF instance and bound to an SPB service. This connects routes within the VRF instance to an SPB service instance (I-SID).

L3 VPN Loopback Interface

Configuring a physical loopback port configuration to define an L3 VPN interface is an available option on OmniSwitch OS6860/6865/6900 platforms. A regular switch port or a static link aggregate can serve as a loopback port. In addition, multiple loopback port pairs are allowed and can be shared between different VRFs.

An L3 VPN loopback interface configuration consists of the following components:

- An IP interface created in a specific VRF instance and bound to a VLAN ID.
- A regular switch port or link aggregate tagged with the IP interface VLAN.
- A service access port that is assigned to an SPB SAP. The SAP encapsulation is configured with the VLAN ID that is tagged on the regular switch port or link aggregate.
- A physical cable that connects the VLAN port with the SAP port to form the loopback configuration.

SPB-M Support on OS9900

The OS9900 now supports SPB-M.

SPB-M Auto-Fabric Support on OS9900

The OS9900 now supports auto-fabric for SPB-M.

VRF "Create" command

To create a new VRF instance the new "create" keyword must now be used.

IPv6 Enhancements on the OS9900

Refer to the feature matrix for a list of supported IPv6 protocols on the OS9900.

OmniVista Cirrus (OmniVista Cloud)

OmniVista Cirrus, the OmniVista Cloud-based solution is an alternative to the current on premise version of OmniVista. The major objectives of a Cloud-based network management system are:

- Reduce the cost of IT for the Customer
- Management from anywhere and ability to manage the network using device ranging from top end workstations to smart phones.
- Simpler network setup
- Easier to use management and monitoring tools.
- Unified wired/wireless management from the cloud.

MACSec

MACSec (MAC Security) provides point-to-point security on Ethernet links between directly connected nodes. MACSec prevents DoS/M-in-M/playback attacks, intrusion, wire-tapping, masquerading, and so on. MACSec can be used to secure most of the traffic on Ethernet links - LLDP frames, LACP frames, DHCP/ARP packets, and so on.

MACSec is supported on the following platforms and ports:

- OS6860(E) – 10G ports on all E/non-E models
- OS6860E-P24Z8/P24 – 1G/10G ports (not supported on 2.5G ports)

This release supports static Key MACSec with support for rotating keys periodically.

1588 transparent clock (e2e only) All Capable Platforms

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.

This release adds support for the OS6860(E)/6865 platforms.

ZeroConf mDNS and SSDP (UPnP/DLNA) Relay Across L3 Networks

Feature support is being enhanced in 8.4.1.R03 by adding the following support for all platforms.

The zero configuration SSDP is extended to responder mode.

The SSDP solution provides the following functionality:

- Allows SSDP enabled devices to discover network services across IP subnet boundaries.

- Allows selective sharing of services based on sharing rules. Service sharing rules can be defined based on VLAN, access role profile(UNP), location, user name and MAC address.

The user name and MAC address attributes are also supported for mDNS.

Four/two-pair Power Over Ethernet

This feature is used to configure 2-pair or 4-pair PoE mode. This command is only applicable for 4-pair high power PoE ports. When 4-pair mode is enabled the switch will consider Alt-A and Alt-B pairs for PoE. When 4-pair mode is disabled the switch will only consider Alt-A pairs for PoE. When 4-pair mode is disabled the maximum PoE power for the port is 30W.

Open Problem Reports and Feature Exceptions

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

System / General / Display

PR	Description	Workaround
226502	Port-mirroring does not work when traffic ingress and egress is on the same GNI-48 and mirror destination is across the VC.	Mirror destination should be on the same Chassis or traffic egress should be a different NI.
229071	The MTU of the internal CMM to NI interface is 9216 on the OS9900. Additional headers need to be added to the IP packets to send the packets from CMM to the NI. The MTU size has been reduced to 9100 to accommodate these headers and is applicable to all NI modules. Any IP packet above max MTU size (9100) will be dropped.	There is no known workaround at this time.
230177	Display error in Egress Packets statistics with "show service counters" command associated with SPB SAP port, counter is NULL.	Use the "show interface counters" command.
230247	On an OS9900 SPB UNP is not supported.	There is no known workaround at this time.
230452	On an OS6860 port monitoring doesn't capture SPB network packets.	There is no known workaround at this time.
230678	When saving RIPng configuration which are configured on a non-default VRF, reloading the switch will delete all the RIPng configurations on non-default VRFs.	There is no known workaround at this time.
230580	On an OS9900 LBD rule is getting overwritten when same ports numbers from different NI participate on LBD.	Avoid using the same port numbers.
231090	The 'show hardware-info' output lists the Microsemi firmware version only for the last module that had the lanpower started.	To list the Microsemi PoE firmware version for modules that support PoE, lanpower needs to be toggled for the specific module whose version needs to be verified in the 'show hardware-info' output.

QoS

PR	Description	Workaround
229438	On an OS6560-P48Z16, a high data traffic rate incoming on ports 1-32 oversubscribing the output on the same range of ports may cause higher priority traffic to be partially dropped over lower priority traffic.	There is no known workaround at this time.

229482	On the OS99-CNI-U8 dhcp-snooping ip-source-filter is not working.	There is no known workaround at this time
229691	Wrong error info for Ip-port with Ip-protocol other than UDP or TCP. "IP protocol can only be 6(TCP) or 17(UDP)" should be "IP protocol can only be 6(TCP) or 17(UDP) with ip-port".	There is no known workaround at this time.
230325	On an OS9900 policy CIR does not work for L2 traffic.	There is no known workaround at this time.

Hardware

PR	Description	Workaround
228368	On OS6900 models, the QSFP-40G-SR-BD transceiver does not come up when used as a VFL link.	There is no known workaround at this time.
228992	The QSFP-4X10G-C5M splitter cable is not supported on the OS6560 switches.	There is no known workaround at this time.
230117	On an OS6865-P16X with a SFP 1G transceiver connected between P16X and P16X/U12X/U28X using 10G ports, the transceiver sometimes gets detected as “SFP_PLUS_COPPER” at one side and as “GBIC_LH,” on the other side, and the link doesn’t come up.	Perform a hotswap (remove/insert) of the transceiver and it will recover.
230313	The QSFP-4X10G-C5M cable is not supported on the OS99-CNI-U8 module.	There is no known workaround at this time.
230555	Multiple link flaps seen when trying to link a 6560-10G interface with the OS9000 NI module OS9-XNI-U12E before the link becomes stable.	There is no known workaround at this time.

Virtual Chassis

PR	Description	Workaround
228412	svcCmm mDREQ & svcCmm mIPMS error messages are seen on console after VC-takeover & mac-flush.	There is no known workaround at this time.

NMS

PR	Description	Workaround
229018	Setting SNMP variable serverPolicyDecision to recachePolicies(1) when using OmniVista Cloud returns	User can ignore error from SNMP operation in a Cloud setup and with a suitable delay check if the policies from LDAP are loaded

	TimeOut error. OmniVista Cloud and associated LDAP resides in the Cloud and due to this change there is an increase in latency delays. The operation to recache continues and if setup was fine, the policies from LDAP are loaded successfully. Users who depend on the status of SNMP Set may have a problem making a determination if the policies were loaded from LDAP or not.	either using CLI command/SNMP/Restful CLI Web-Interface. Please do not retry SNMP set on serverPolicyDecision in case of timeout failure - it will only make the problem worse by initiating the policy recache operation one more time.
230529	Trying to load a PolicyList from LDAP is failing on an OS6560. Policy objects other than PolicyList are loaded successfully from LDAP.	Create the PolicyList/Policies using CLI on an OS6560.

MACSec

PR	Description	Workaround
231085	When the command 'clear interfaces port <port_num> macsec-statistics' is executed the port gets reconfigured for dynamic mode. This configuration is retained on the display even after the configuration is removed.	Reconfigure the MACSec mode to static.

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-XNI-P48Z16	OS99-XNI-P48Z16
OS99-CNI-U8	OS99-CNI-U8

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.R03.

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	
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 – IPv6	Y	Y	Y	N	Y	IPv6 multicast routing and MLDv1 and v2 are not supported in VRF
VRF – DHCP Client	Y	Y	Y	N	Y	
Automatic/Intelligent Fabric	Y	Y	Y	Y	Y	
Automatic VC	Y	Y	Y	Y	N	
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	N	Y	
Remote Chassis Detection (RCD)	Y	N	N	N	Y	
Web Services & CLI Scripting	Y	Y	Y	Y	Y	

Feature	OS6900	OS6860(E)	OS6865	OS6560	OS9900	Notes
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	
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	
BGP route reflector for IPv6	Y	Y	Y	N	Y	
BGP ASPATH Filtering for IPv6 routes on IPv6 peering	Y	Y	Y	N	Y	
BGP support of MD5 password for IPv6	Y	Y	Y	N	Y	
BGP 4-Octet ASN Support	Y	Y	Y	N	Y	
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	
IPv6 DHCP relay and Neighbor discovery proxy	Y	Y	Y	N	Y	
ISIS IPv4/IPv6	Y	Y	Y	N	N	
M-ISIS	Y	Y	Y	N	N	
OSPFv3	Y	Y	Y	N	Y	
RIP v1/v2	Y	Y	Y	Y	Y	
RIPng	Y	Y	Y	N	Y	
DHCP Server (v4, v6 with integrated support of QIP remote management)	Y	Y	Y	Y	Y	IPv6 server not supported on OS6560.
VRRP v2	Y	Y	Y	Y	Y	
VRRP v3	Y	Y	Y	N	Y	
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	
DHCPv6 Relay	Y	Y	Y	N	Y	
IP Multinetting	Y	Y	Y	Y	Y	
IPSec	Y	Y	Y	N	N	
Server Load Balancing (SLB)	Y	Y	Y	N	N	

Feature	OS6900	OS6860(E)	OS6865	OS6560	OS9900	Notes
Multicast Features						
IGMP v1/v2/v3	Y	Y	Y	Y	Y	
IPv4 Multicast Switching	Y	Y	Y	Y	Y	
DVMRP	Y	Y	Y	N	Y	
IPv6 Multicast Switching (MLD v1/v2)	Y	Y	Y	N	Y	
IPv6 Scoped Multicast Addresses	Y	Y	Y	N	N	
PIM-DM	Y	Y	Y	N	Y	
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	
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	N	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	Y	PVST+ not supported on OS6560/9900
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	N	N	
ERP v1/v2	Y	Y	Y	N	N	
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	
Private VLANs	Y	Y	Y	N	N	

Feature	OS6900	OS6860(E)	OS6865	OS6560	OS9900	Notes
SIP Snooping	N	Y	N	N	N	
SPB	Y	Y	Y	N	Y	
VXLAN	Q32/X72	N	N	N	N	
VM/VXLAN Snooping	Y	N	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	N	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	N	N	
Ingress/Egress bandwidth limit	Y	Y	Y	Y	Y	
Tri-color marking	Y	Y	Y	N	N	
QSP Profiles 2/3/4	Y	Y	Y	N	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	
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	
LLDP	Y	Y	Y	Y	Y	
TACACS+ Client	Y	Y	Y	Y	Y	
TACACS+ command based authorization	Y	Y	Y	N	Y	

Feature	OS6900	OS6860(E)	OS6865	OS6560	OS9900	Notes
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	N	N	
Radius test tool	Y	Y	Y	Y	Y	
Storm Control	Y	Y	Y	N	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	
EVB	Y	N	N	N	N	
FCoE / FC Gateway	Y	N	N	N	N	
FIP 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.R03 (GA)**
OS6900	8.3.1.160.R02 (GA) 8.3.1.180.R02 (MR) 8.4.1.170.R01 (GA) 8.4.1.229.R02 (GA) 8.4.1.233.R02 (MR)
OS6860(E)	8.3.1.160.R02 (GA) 8.3.1.179.R02 (MR) 8.3.1.180.R02 (MR) 8.4.1.170.R01 (GA) 8.4.1.229.R02 (GA) 8.4.1.233.R02 (MR)
OS6865	8.3.1.160.R02 (GA) 8.3.1.180.R02 (MR) 8.4.1.170.R01 (GA) 8.4.1.229.R02 (GA)
OS6560	8.4.1.229.R02 (GA)
OS9900	8.4.1.229.R02 (GA) **Note**: ISSU supported on VC of 2 only.
Note: For any switch with a multicast configuration ISSU is only supported from 8.4.1.R02 GA or MR. Earlier releases must use a standard upgrade.	

8.4.1.R03 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 8.4.1.233.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 – Uos.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           8.4.1.141.R03            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 – Uos.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           8.4.1.141.R03            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.R03.

PR	Summary
230299	OS9900: Port scan shows port# 52655 as unknown.
230297	VC of seven OS6860E crashed due to lldpcmm task and generated the PMD file.
229965	OS6900: Switch rebooted couple of times due to "Segmentation Fault-lldpNi"
229892	OS9900:CMM has different PS LED status at the Secondary CMM module.
229866	Error message: "lldpCmm library(plApi) error(2) plGetIfIndexFromBasePort@1649" repeating in the swlo
229865	LLDP display VF-links after multiple the third vc-takeover
229828	Route leaking issues after upgrade from 841.R01 to 842.229.R01 (GA)
229766	Randomly SNMP polling has error on MIB dot3adAggActorAdminKey
229756	an 6860 , SPB node, running, 831.314.R01, with loopback (linkagg 31) and sap linkagg 20 (connected
229753	NOKIA-SAM: ipmscmm application crashes upon receiving SNMP get requests for non-existent vlan/servic
229663	OSPFv3 neighborship takes 4-6 mn to establish once a failover happens on firewalls (master/backup) c
229602	OS9900: tcamCmm app warning messages seen on the swlogs
229541	OS6860 crash with the PMD file.
229474	at bootup, L2profile for SPB tried to be applied before qosNi is ready to receive messages from qosC
229399	OS6860: No mac-addresses learnt on the slave chassis ports shown for the object "slMacAddressGblRowS
229369	LGS 266: potential resource leak forking CLI message
229364	LGS 262: Memory not freed parsing ipv6 ping and traceroute commands
229363	LGS 261: Memory not freed parsing ipv4 traceroute and ping commands
229360	LGS 258: Getting CLI prompt can overwrite buffer (CLI lib)
229356	LGS 254: String copy of configuration file name not NULL terminated
229354	LGS 253: Resource leak clearing error files
229295	OS6860: Though authentication is successful, UNP user goes to LPS Blocking status without any violat
229278	Default route is not learned in OSPF on VC-takeover, when graceful restart is enabled
229128	Unable to change the password for the admin account
229051	OS6900- Unable to allow local user to manage LBD, when TACACAS is unavailable
228946	BGP neighbors up time is more than 100 years.
228921	Wrong behavior when Policy Rule is applied from OV and stats issue on show active policy
228896	OS-XNI-T8 port Auto-negotiation shows "10-F" though it is not supported
228892	SFP-GIG-SX - Accepts 40000 M speed
228890	SFP-10G-LR accepts 40000 and 1000M in the configuration, though it is not updated on the interface
228850	OS6860(8.4.1.170.R01): UNP user stays in LPS Blocking status and no issue in traffic flows

228818	LGS 251: File handler resource leak
228817	LGS 250: Buffer not NULL terminated when copying CLI command
228816	LGS 249: SNMP resource memory leak
228814	LGS 248: Use after free processing SNMP OIDs
228801	LGS 235: Reflective cross-site scripting error retrieving IP interface statistics
228789	LGS 231: Use of uninitialized pointer reading in config file in conf_util.c
228766	SPB switch with routing configuration fails to ping a L2 switch connected by a static linkagg.
228682	9900: tcam errors configuring policy validity period combo
228500	"Pse Port Detection" filling up the logs
228434	OS9900: High CPU on NI-3 due to portmgrni task.
227690	SPB service access port loosing connectivity to the edge device.
226591	LGS 124:Out of bounds memory access in performance tools
226554	3XOS6900: FCOE is not working for a new vlan
224353	OS6860: QoS policy on 6860E applied from OV is not working

Appendix F: PoE Firmware Upgrade Instructions

Follow the instructions below to upgrade the PoE firmware. The upgrade must be done one module at a time. This procedure requires the lanpower service to be stopped and restarted which will cause any connected powered devices to lose power.

Step 1. Download the PoE Firmware from the Service & Support Site

Visit **support.esd.alcatel-lucent.com** and click on **Download Software and/or Product Documentation -> OmniSwitch 9900** to download the PoE firmware. The file name is "14026504_0621_023.s19".

Step 2. FTP the PoE Firmware File to the OS9900

FTP the PoE firmware file to the **/flash** directory of the OS9900.

```
C:\temp\poe_firmware>ftp 10.255.13.45
Connected to 10.255.13.45.
220 (vsFTPd 3.0.2)
User (10.255.13.45:(none)): admin
331 Please specify the password.
Password:
230 Login successful.
ftp> bin
200 Switching to Binary mode.
ftp> hash
Hash mark printing On ftp: (2048 bytes/hash mark).
ftp> pwd
257 "/flash"
ftp> put 14026504_0621_023.s19
200 PORT command successful. Consider using PASV.
150 Ok to send data.
################################################################################
226 Transfer complete.
ftp: 165043 bytes sent in 0.02Seconds 8686.47Kbytes/sec.
ftp> quit
221 Goodbye.
```

Step 3. Determine Current PoE Firmware Version

The **show hardware-info** command will only show the PoE firmware version for the last module on which the lanpower service was started (See PR 231090). To check the PoE version for a specific module use the **lanpower slot service** command to restart the lanpower service.

```
-> lanpower slot 1/4 service stop
-> lanpower slot 1/4 service start
-> show hardware-info
Chassis 1
CPU Manufacturer                  : GenuineIntel
CPU Model                         : Intel(R) Atom(TM) CPU  C2518  @ 1.74GHz
Compact Flash Manufacturer        : USB
Compact Flash size                : 1964617728 bytes
RAM Manufacturer                  : Other
RAM size                          : 16322936kB
Control FPGA version              : 2.3.0
Power FPGA version                : 0.8
Coreboot Version                  : 8.4.1.103.R01
CFMs Present                      : 1,2
Power Supplies Present            : 1,2,3
Fan Trays Present                 : 1,2,3
```

```
NIs Present                       : 1,2,3,4,5,6,7

POE:  PD69100 Software Version 00.0263.01 Hardware Version 00 NI 4
```

Step 4. Upgrade PoE Firmware for the Module

The lanpower service must be stopped on ALL PoE modules before upgrading the PoE firmware on any module. If the lanpower service is running on any module in the chassis while trying to perform the upgrade an error similar to “ERROR: Chassis 1 Active” will be displayed. It can take between 3 and 4 minutes for the upgrade to complete for each module.

```
-> lanpower slot 1/3 service stop
-> lanpower slot 1/4 service stop
-> lanpower slot 1/4 update-from 14026504_0621_023.s19

Thu Jan 11 07:08:29 : lpCmm LanCmm error message:
+++ lpCmmUpdateFileProgramNI: 344
+++ lpCmmUpdateFileProgramNI: 416 chassisId 4 slot 1
+++ lpCmmUpdateFileProgramNI: 423

Thu Jan 11 07:08:29 : lpCmm LanCmm info message:
+++ Reprogramming Sequence Started  2217 chassisId 1 slot 4

Thu Jan 11 07:08:39 : lpCmm LanCmm info message:
+++ Controller Memory Sequence Begining 623 chassisId 1 slot 4

Thu Jan 11 07:08:46 : lpCmm LanCmm info message:
+++ Controller Memory Reflash Please Wait 465 chassisId 1 slot 4

Thu Jan 11 07:12:04 : lpCmm LanCmm info message:
+++ Reprogram Pass 2221 chassisId 1 slot 4
```

Step 5. Restart Lanpower for the Module After Upgrade

To confirm a successful upgrade, restart the lanpower service for the upgraded module and use the **show hardware-info** command to verify the PoE firmware version.

```
-> lanpower slot 1/4 service start

-> show hardware-info
Chassis 1
CPU Manufacturer                  : GenuineIntel
CPU Model                         : Intel(R) Atom(TM) CPU  C2518  @ 1.74GHz
Compact Flash Manufacturer        : MICRON
Compact Flash size                : 1997967360 bytes
RAM Manufacturer                  : Other
RAM size                          : 16323004kB
Control FPGA version              : 2.0.0
Power FPGA version                : 0.8
Coreboot Version                  : 8.3.1.103.R01
CFMs Present                      : 1,-
Power Supplies Present            : 1,2,-
Fan Trays Present                 : 1,2,3
NIs Present                       : 1,-,3,4,-,-,-

POE:  PD69100 Software Version 00.0265.04 Hardware Version 00 NI 4
```

Step 6. Repeat the Upgrade Procedure as Needed

Repeat steps 3 through 5 for any additional modules that need to be upgraded. Once all modules have been upgraded be sure to restart the lanpower service for all modules.