Release Notes

OmniSwitch 6900/10K

Release 7.3.4.R01

These release notes accompany release 7.3.4.R01 software which is supported on the OmniSwitch 6900 and OmniSwitch 10K platforms. These release notes provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

[IMPORTANT] *MUST READ* - This release changes default AOS behavior as well as deprecating some feature support. It is required that the PREREQUISITE section be read and UNDERSTOOD prior to upgrading to AOS Release 7.3.4.R01. If, after reading the PREREQUISITE section, you still have questions, please contact Service & Support for further clarification.

Contents

Related Documentation

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

OmniSwitch 6900 Series Hardware User Guide

Complete technical specifications and procedures for all OmniSwitch Series chassis, power supplies, and fans.

OmniSwitch 10K Getting Started Guide

Describes the hardware and software procedures for getting an OmniSwitch up and running.

OmniSwitch 10K Hardware User Guide

Complete technical specifications and procedures for all OmniSwitch Series chassis, power supplies, and fans.

OmniSwitch AOS Release 7 CLI Reference Guide

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

OmniSwitch AOS Release 7 Network Configuration Guide

Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethernet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access), Quality of Service (QoS), and link aggregation.

OmniSwitch AOS Release 7 Switch Management Guide

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

OmniSwitch AOS Release 7 Advanced Routing Configuration Guide

Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, OSPFv3, and IS-IS.

OmniSwitch AOS Release 7 Data Center Switching Guide

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

OmniSwitch AOS Release 7 Transceivers Guide

Includes SFP, SFP+, and QSFP transceiver specifications and product compatibility information.

Technical Tips, Field Notices

Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

OmniSwitch 6900 Series Release 7.3.4.R01 requires 2GB (6900X models) / 4GB (6900T models) / 8GB (OS6900-Q32) of SDRAM and 2GB flash memory. This is the standard configuration shipped.

OmniSwitch 10K Series Release 7.3.4.R01 requires 4GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.

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

UBoot and FPGA Requirements

The software versions listed below are the MINIMUM required, except where otherwise noted. Switches running the minimum versions, as listed below, do not require any UBoot or FPGA upgrades. 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 7.3.4.R01 AOS software available from Service & Support.

- A separate file containing the Uboot and FPGA upgrade files is available from Service & Support.
- Please refer to the Upgrade Instructions section at the end of these Release Notes for step-by-step instructions on upgrading your switch to support 7.3.4.R01.

OmniSwitch 6900-X20/X40 – AOS Release 7.3.4.450.R01(GA)

Hardware	Uboot	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[1]
All Expansion Modules	N/A	N/A

1. FPGA 1.3.0/2.2.0 is required to support the XNI-U12E (Introduced in 7.3.3.R01)

OmniSwitch 6900-T20/T40 – AOS Release 7.3.4.450.R01(GA)

Hardware	Uboot	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[1]
All Expansion Modules	N/A	N/A[2]

1. FPGA 1.6.0 is required to support the XNI-U12E (Introduced in 7.3.3.R01)

OmniSwitch 6900-Q32 – AOS Release 7.3.4.450.R01(GA – Factory Default)[2]

Hardware	Uboot	FPGA
CMM	7.3.4.277.R01[1]	0.1.8/0.0.0[1]
All Expansion Modules	N/A	N/A

1. Shipped from factory with correct version, no upgrade is available or required.
2. AOS 7.3.4.R01 is the minimum version supported. The OS6900-Q32 cannot be downgraded.

OmniSwitch 10K - Release 7.3.4.450.R01 (GA)

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

[IMPORTANT] *MUST READ*: AOS Release 7.3.4.R01 Prerequisites

Please note the following important release specific information prior to upgrading or deploying this release. The information below covers important upgrade requirements, changes in AOS default behavior, and the deprecation of features.

- Prior to upgrading to AOS Release 7.3.4.R01 please refer to Appendix A for important best practices, prerequisites, and step-by-step instructions.
- All switches that ship from the factory with AOS Release 7.3.4.R01 will default to VC mode and attempt to run the automatic VC, automatic remote configuration, and automatic fabric protocols. Please refer to Auto Management Features for additional information. All OS6900-Q32s will ship from the factory with AOS Release 7.3.4.R01.
- Multi-Chassis Link Aggregation is no longer being supported in AOS Release 7.3.4.R01. All CLI and MC-LAG functionality has been deprecated. If Multi-Chassis Link Aggregation support is required DO NOT upgrade to AOS Release 7.3.4.R01 and contact Service & Support.
- If upgrading from AOS Release 7.3.1 note the following:
    - VRF functionality was updated to use the new profiles capability in 7.3.2.R01. These new profiles are not compatabile with earlier versions of AOS. It's strongly recommened to create a backup of the 7.3.1 configuration prior to upgrading to prevent the VRF configuration having to be rebuilt if a switch should need to be downgraded.
    - A new predefined DCB profile 11 was introduced in 7.3.2.R01, this will overwrite any existing custom profile 11.

New Hardware Support

OS6900-Q32

40-Gigabit Ethernet fixed conguration chassis in a 1U form factor with thirty-two (32) 40GBase-X QSFP+ ports, redundant AC or DC power and front to back cooling. The switch includes:

- 1 - Console Port (USB Form Factor - RS-232)
- 1 - USB Port (For use with Alcatel-Lucent OS-USB-FLASHDR USB flash drive)
- 1 - EMP Port
- 32 –QSFP+ Ports
- 1 Slot – Fan Tray
- 2 Slots – Power Supplies (AC or DC)

Port Groups 1 – 6, 11-22, and 27-32 support 4X10G splitter cables for up to ninety-six (96) 10-Gigabit ports. When a splitter cable is used the port numbering scheme changes to accommodate the 4 10-Gig ports by using letters a, b, c, d to refer to the 10-Gig sub-ports. When referring to a single sub-port the port letter should be used to differentiate between all the sub-ports. If no letter is given the command assumes port 'a', for example.

```
-> show interfaces 1/1/1 – refers to interface 1/1/1a
-> show interfaces 1/1/1a – refers to interface 1/1/1a
-> show interfaces 1/1/1d – refers to interface 1/1/1d
```

When referring to a range of ports the lettered sub-ports are implied, for example:

```
-> show interfaces 1/1/1-2 – refers to interfaces 1/1/1a, 1b, 1c, 1d and 1/1/2a, 2b, 2c, 2d
-> show interfaces 1/1/1a-1c – refers to interfaces 1/1/1a, 1b, 1c
-> show interfaces 1/1/1-2a – refers to interfaces 1/1/1a, 1b, 1c, 1d, and 1/1/2a.
```

SFP-10G-ZR Transceiver

10-Gigabit optical transceiver (SFP+). Supports data transmission at 1550nm over up to 80km single mode fiber. LC connector type.

QSFP-4X10G-SR Transceiver

Four channel 40-Gigabit QSFP+ Optical Transceiver with MPO connector. Can be used as a 40GBase-SR 40-Gigabit connection or as a breakout to four(4) 10G-Base-SR 10-Gigabit connections.

QSFP-4X10G-C Transceiver

Four channel 40-Gigabit QSFP+ Direct Attached Copper Splitter Cable. Connects a single 40G QSFP+ port to four(4) 10G SFP+ ports. Available in 1/3/5 meter lengths.

New Software Features and Enhancements

The following software features are being introduced with the 7.3.4.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

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

7.3.4.R01 New Feature/Enhancements Summary

Feature	Platform	License
Data Center Feature Support		
- VXLAN	6900-Q32	Data Center
- VM/VXLAN Snooping	6900/10K	Data Center
DHCP		
- Internal DHCPv4 and DHCPv6 Server	6900/10K	Base
- IPv6 DHCP Relay Agent	6900/10K	Base
- DHCP Snooping	6900/10K	Base
Layer 3 Feature Support		
- BGP 4-Octet ASN	6900/10K	Advanced
- BGP AS Path Filtering for IPv6	6900/10K	Advanced
- BGP Password Support for IPv6	6900/10K	Advanced
- BGP Route Reflector for IPv6	6900/10K	Advanced
- Distributed ARP	6900	Base
- Increase OSPFv2 Interfaces	6900/10K	Advanced
- ISIS for IPv6	6900/10K	Advanced
- M-ISIS	6900/10K	Advanced
- Static Routing to an IP Interface Name	6900/10K	Base
- IP Routed Port	6900/10K	Base
Automatic Management Feature Support		
- Automatic Virtual Chassis	6900/10K	Advanced
- Automatic Remote Configuration	6900/10K	Base
- Automatic Fabric	6900/10K	Base
- Automatic IP Protocols	6900/10K	Base

Feature	Platform	License
Management Feature Support		
- Embedded Pyton Scripting / Event Manager Support	6900/10K	Base
- IP Managed Services	6900/10K	Base
- OpenFlow Support for Standalone and Virtual Chassis	6900/10K	Base
Security		
- 802.1x for VLANs, SPBM, and VXLAN Services	6900/10K	Base

Data Center Feature Descriptions

VXLAN

A Virtual eXtensible Local Area Network (VXLAN) is a Layer 2 overlay network that is used to segment and tunnel Virtual Machine (VM) traffic through a data center or cloud network infrastructure. The OmniSwitch implementation of this feature introduces the following benefits into the network:

- Provides a gateway device that sits at the edge of a VXLAN UDP tunnel to serve as a VXLAN Tunnel End Point (VTEP). An OmniSwitch VXLAN gateway offers a high bandwidth, low latency option for connecting VM traffic to remote servers or other VMs over a Layer 3 network.
- Provides Layer 2 connectivity between devices in the same bridging domain over an IP transport network. For example, a Virtual Machine (VM) can communicate across a Layer 3 network with a remote VM as long as both VMs reside in the same VLAN domain on either side of the Layer 3 network.
- Increases the scalability of the network beyond the limit of 4096 VLANs. A VXLAN Network Identifier (VNI) is used to isolate VLAN traffic into logical network segments. Up to 16 million logical networks (segment IDs) are possible when VXLAN is implemented.
- Conserves MAC address table space by allowing duplicate MAC addresses to reside within the VXLAN domain, as long as each address is associated with a different VNI.
- Transparently extends the Layer 2 network by connecting VLANs from multiple hosts through VXLAN (UDP) tunnels.
- Provides Layer 2 migration of a Virtual Machine (VM) across a Layer 3 infrastructure to a remote server host; without VXLAN, Layer 2 migration is restricted to other servers within the local Layer 2 network.

.

VXLAN/VM Snooping

The OmniSwitch Virtual Machine (VM) Snooping feature attempts to detect and identify Virtual eXtensible LAN (VXLAN) traffic by inspecting packets to determine if they are VXLAN encapsulated packets. Once VXLAN traffic is identified, VM Snooping collects and stores information about the flow in a database on the local switch. Additional configurable options for this feature include the ability to apply QoS policy list rules to the identified flow and generating SNMP traps when a new VM is learned.

Using this implementation of VM Snooping, an administrator can determine the switch interface on which VXLAN packets are entering the network along with other details about the VM traffic.

- Provide visibility of the VNIs and the virtual machines within each VNI
- Allow QoS policies to be dynamically applied to VXLAN frames based on header fields and the fields of the internal packet.
- Per port enable/disable
- Database of discovered VMs/VNIs, with statistics
- Two filtering modes to fit needed granularity and scalability
- QoS policies offer wide range of control
- Traps are available for VM learning, aging out, and resource limits
- Can be used on VNP/UNP ports (packets otherwise sent to the CPU)

DHCP Feature Descriptions

Internal DHCPv4 and DHCPv6 Server

The OmniSwitch now supports an internal DHCP Server compliant with RFC 2131 and RFC 3315 based on Vital QIP 8.0 release. This feature can be used to provide IP addresses for small offices, management network, or local phone services. The following files are used to configure the internal DHCP server setting on the OmniSwitch:

- IPv4 Policy Files- dhcpd.conf, dhcpd.pcy
- IPv6 Configuration Files – dhcpd6.conf, dhcpd6.pcy

DHCP Policy files - The dhcpd(v6).pcy files initialize the global attributes for the DHCP server.

DHCP Configuration files - The dhcpd(v6).conf files are used to configure specific DHCP server settings on the switch such as the following:

- MAC pool allowed (for DHCPv4)
- MAC pool excluded (for DHCPv4)
- Subnet pools
- Dynamic scopes
- Static scopes
- IP range, mask, DNS, Default router, NetBIOS configurations for DHCPv4.
- User class specific configs.
- Vendor class specific configs.
- DUID Pool (for DHCPv6 only).
- Excluded DUID Pool (for DHCPv6 only)
- Manual DUID mapping (for DHCPv6 only).
- Other options that need to be sent to the client.

DHCP Relay Agent for IPv6

The Alcatel-Lucent OmniSwitch implementation of RFC 3315 contains IPv6 support. A DHCPv6 relay agent is required in situations where DHCPv6 clients do not reside on the same link as the DHCP server. The DHCPv6 Relay on OmniSwitch processes and forwards all DHCPv6 messages triggered by DHCPv6 client to the configured DHCPv6 relay agent as a unicast packet.

DHCPv6 Relay is a per-interface option that can be enabled on any IPv6 interface.

The DHCPv6 relay can operate only in the default VRF. It supports multicast-capable IPv6 interfaces (VLAN and configured tunnel interfaces) and non-multicast-capable IPv6 interfaces (6to4 tunnel).

The DHCPv6 Relay agent is part of the link-scoped multicast group (FF02::1:2) on the interface. Any messages sent by a client to that address will then be handled by DHCPv6 Relay agent.

A maximum of five unicast or link-scoped multicast relay destinations can be configured for each interface on which DHCPv6 Relay is enabled. The DHCPv6 relay for the interface will be automatically disabled when all the relay destinations configured for that interface are removed.

DHCP Snooping

DHCP snooping provides network security by filtering untrusted DHCP messages by building and maintaining a DHCP snooping binding database. It acts like a firewall between untrusted hosts and DHCP servers. This feature prevents the normal flooding of DHCP Discover/Request and DHCP Offer packets. These packets will instead be delivered only to the appropriate DHCP server and client ports respectively.

Ingress Source Filtering

In addition to filtering untrusted DHCP messages, DHCP Snooping allows user to configure Ingress Source Filtering as a security feature.

When Ingress Source Filtering (ISF) is enabled on a port or linkagg port, the initial packets permitted for traffic are DHCP, DNS and ARP, so as to allow the client to obtain and IP address from the DHCP server in which a MAC-IP Binding entry is created in the DHCP Snooping task, it will then allow packets that match the IP address/MAC address/ port combination that is obtained from the DHCP snooping binding table entry. Other non-matching packets will be dropped.

When ISF is enabled on a VLAN then the VLAN ID is added to the matching criteria as an additional parameter that must be matched.

DHCP Relay Agent Information Option-82

The DHCP Option-82 feature enables the relay agent to insert identifying information into client-originated DHCP packets before the packets are forwarded to the DHCP server. The implementation of this feature is based on the functionality defined in RFC 3046.

When DHCP Option-82 is enabled, communications between a DHCP client and a DHCP server are authenticated by the relay agent . To accomplish this task, the agent adds Option-82 data to the end of the options field in DHCP packets sent from a client to a DHCP server.

User-configurable Option 82 Suboption Format — Allows the user to specify the type of information (switch base MAC address, system name, or user-defined string) that is inserted into the Circuit ID and Remote ID suboptions of the Option-82 field. This functionality only applies when DHCP Snooping Option-82 Data Insertion is enabled.

Layer 2 Feature Descriptions

SPB Loopback Detection

A provider network with a set of multiple switches interconnected together can be logically viewed as a large single switch. The large single switch provides service access points to customers' networks. Configuration faults in customer networks can result in loops spanning both provider and customer networks. This can result in broadcast storms. In order to protect provider's network from broadcast storms, loops that involve SAP ports need to be detected and broken.

The LBD feature can detect and break loops created on a service-access interface. For a service-access interface the LBD feature can be enabled for a specific port or linkagg. LBD for service-access points allows shutting down only the specific interface of the link involved in the loop.

Layer 3 Feature Descriptions

BGP 4-Octet Autonomous System Number (ASN)

This feature enhancement provides the following:

- BGP Support for 4-octet (32 bit) ASN for BGP neighbor interoperability and path attribute interoperability as per RFC 6793.
- Capabilities Advertisement with BGP-4 - The advertisement and discovery of 4-octet ASN capability by using the BGP CAPABILITY FIELDS.
- Support for two new optional transitive attributes AS4_PATH and AS4_AGGREGATE. These attribute are used while interacting with NEW BGP speaker and OLD BGP speaker.
- To establish a neighbor relationship between non-mappable BGP 4-octet ASNs with BGP 2-octet ASNs the reserved 2-octet ASN AS_TRANS 23456 is used.
- Extended Community will be used for non-mappable 4-octet ASNs with BGP 2-octet ASNs.
- The 4-octet ASN is represented in one of three ways:
    - asplain (simple decimal notation)
    - asdot+ (two 16-bit values as low-order and high-order)
    - asdot (a mixture of asplain and asdot+).

BGP AS Path Filtering for IPv6

BGP selects routes for subsequent advertisement by applying policies available in a pre-configured local Policy Information database. This support of policy-based routing provides flexibility by applying policies based on the path (that is, AS path list), community attributes (that is, community lists), specific destinations (that is, prefix lists and prefix6 lists) and so on. You could also configure route maps to include all of the above in a single policy. For BGP to do policy-based routing, each BGP peer needs to be tied to an inbound and/or outbound policy. Each one of the above policies can be assigned as an in-bound or out-bound policy for a peer. AOS BGP currently supports the following policies:

- AS Paths. An AS path list notes all ASs that the route travels to reach its destination.
- Community List. Communities can affect route behavior based on the definition of the community.
- Prefix List. Prefix list policies filter IPv4 routes based on a specific IPv4 network address, or a range of IPv4 network addresses.

- Prefix6 List. Prefix6 list policies filter IPv6 routes based on a specific IPv6 network address, or a range of IPv6 network addresses.
- Route Map. Route map policies filter routes by amalgamating other policies into one policy.

All of these policies are currently supported for both IPv4 and IPv6 routes/peers. The route-map configuration supports IPv4 and IPv6 prefix lists as well as IPv6 prefixes and masks.

BGP MD5 Authentication Support for IPv6

Similar to IPv4 neighbors, an IPv6 neighbor can now be configured to use MD5 authentication for the IPv6 TCP connections. This sets an encrypted MD5 signature for TCP sessions with this IPv6 peer in compliance with RFC 2385. If the password mentioned and IPv6 address of the peer do not match, then the local router will not communicate with that peer.

BGP Route Reflector for IPv6

Similar to IPv4 neighbors, an IPv6 neighbor can now be configured as client for route reflector.

The BGP specification states that a BGP speaker cannot advertise a route to an I-BGP neighbor if that BGP speaker originally heard the route from another I-BGP speaker. This results in a requirement for a full mesh of I-BGP sessions within an AS. When route reflection is configured on internal BGP speakers in an autonomous system, the topology does need not be fully meshed. The route reflector takes responsibility for passing internal BGP-learned routes to its peers.

Distributed ARP

The distributed ARP feature enables a greater number of hardware ARP lookups. This is accomplished by distributing ARPs across the individual chassis comprising an OS6900 virtual chassis. The feature dynamically designates a specific Network Interface (NI) as the designated-NI for all ARP entries on an IP interface basis. All packets destined for the host in ARP lookup table will be routed to the designated NI for routing.

Designated-NI and Distributed ARP Management

The designated-NI is dynamically assigned for the IP interface. By default the NI with the most number of active ports in the VLAN is set as the designated-NI. If the underlying VLAN has a linkagg with members on multiple chassis, one NI on each chassis will be chosen as the designated-NI.

When the total number of ARPs learned on the designated-NI exceed a fixed percentage of capacity (e.g. 95%), a new designated-NI is chosen for the IP interface. The NI with the second highest active ports in the VLAN is selected (provided space is available on the new NI).

All packets destined for the IP interface will be routed to the designated NI for lookups. The designated-NI performs the ARP lookup, and forwards the traffic. If the designated-NI does not have a matching ARP entry, the traffic is trapped to the CPU of that NI, which will then resolve the ARP.

Using the designated NI method, there is a possibility of packets traversing the backplane/VFL multiple times. In cases where backplane throughput is a concern configure the network in such a way as to minimize the backplane utilization while optimally utilizing hardware (i.e. centralizing VLAN ports to NIs).

Increase OSPFv2 Interfaces

This enhancement helps increase the number of supported OSPF interfaces per Area. In this enhancement, the "**areamaxintfs**" variable can be set to the required number of OSPF interfaces using the "**debug ip ospf set areamaxintfs <num>**" command . An optimization has also been introduced with this enhancement.

Before this enhancement, a passive OSPF interface was created with four lines in the configuration file. Now a route map is used. A route map with set action of route-type internal can be created for the local interface (routes) on which a passive OSPF interface needs to be created. Using this route-map in redistribution of local into OSPF, the passive OSPF interfaces will be learned as intra routes. Thus those interfaces will act as passive OSPF interfaces. The OSPF interfaces created by the route-map command can be accessed in all the OSPF display commands. This passive OSPF interface will not be written into boot.cfg and will not be visible in snapshot.

Example:

Include all IP interfaces which need to be configured as passive OSPF interface, in a route map and then use the below commands to have them as passive OSPF interface without configuring those IP interfaces as OSFP interfaces.

-> ip route-map "R1" sequence-number 50 action permit

-> ip route-map "R1" sequence-number 50 set metric-type internal

-> ip redist local into ospf route-map R1 status enable.

ISIS for IPv6

Intermediate System-Intermediate System (IS-IS) is a shortest path first (SPF) or link-state protocol. IS-IS is an interior gateway protocol (IGP) that distributes routing information between routers in a single autonomous system (AS) for IP (IPv4 and IPv6) as well as OSI environments. This feature allows a single routing protocol to support pure IP and OSI environments, and dual environments. Integrated IS-IS is also deployed extensively in an IP-only environment.

M-ISIS

Multi-topology (M-ISIS) support is necessary in IS-IS to support network domains in which non-dual stack IS-IS routers exist. The default protocol behavior of IS-IS is to construct shortest paths through the network using the routers' MAC addresses with no regard to the different IP address families supported. This behavior may result in black-holed routing when there are some IPv4-only or IPv6-only routers in an IS-IS routing domain, instead of all dual-stack routers. M-ISIS mechanism runs multiple, independent IP topologies within a single IS-IS network domain, using separate topology-specific SPF computation and multiple Routing Information Bases (RIBs). M-ISIS is advised in networks containing ISIS enabled routers with a combination of IPv4 and IPv6 capabilities.

Static Routing to an IP Interface Name

The feature allows using the interface name instead of the specific IP address and index to configure static route and static ARP.

IP Routed Port

This features allows for the configuration of a IP interface that is associated to a physical port or link aggregate and VLAN on which layer 3 functionality is supported. The underlying routed VLAN does not support layer 2 functionality. The routed VLAN allows for the configuration of the physical port, IP interface and VLAN in a single command.

Automatic Management Feature Descriptions

All switches that ship from the factory with 7.3.4.R01 will default to Virtual Chassis mode and attempt to run the automatic VC, automatic remote configuration, and automatic fabric protocols. The features can be disabled during the switch reboot or after the switch has finished booting if desired. The use of these features

will require either the Demo or Advanced license. A 45-day demo is included for the Advanced license on switches shipped from the factory with AOS Release 7.3.4.

When a switch boots with no configuration file or with a configuration file with a size of 0 bytes, the following boot processes occur:

- The switch will run the automatic VC protocol and attempt to automatically configure the VFLs and setup a VC. This process can take approximately 10-20 minutes to complete. Please refer to the ***Getting Started*** chapter of the Switch Management Guide for additional information.
- Once the automatic VC process is complete, the automatic remote configuration process will begin. It can take approximately 180 seconds for this process to complete if there is no remote configuration server available. Refer to the ***Getting Started*** chapter of the Switch Management Guide for additional information.
    - The automatic remote configuration download process can be aborted at any time by entering **auto-config-abort**.
- Once the automatic remote configuration process completes, the automatic fabric process will begin. Refer to the ***Getting Started*** chapter of the Switch Management Guide for more information.
    - The automatic fabric process can be disabled at any time by entering '**auto-fabric admin-state disable'**.

This boot process only applies to switches that boot without a configuration file, such as newly shipped switches from the factory.

The automatic management features can be disabled at the start of the switch boot process by pressing '**Y**' when prompted. This disables all the automatic management features and the switch will come up in standlone mode.

To prevent a switch from re-running the automatic fabric process upon the next reboot enter '**write memory**' to save the configuration to the configuration file.

If the automatic management features are not disabled while the switch is booting the **auto-fabric admin-state disable remove-vc-reload** command can be entered. This will clear any automatic configuration and reboot the switch into standlaone mode.

Automatic Virtual Chassis

Automatic Virtual Chassis can be used to ease the required manual configuration for a VC. The automatic VC feature will allow a brand new chassis shipped from the factory or a chassis with no configuration to be setup as a VC without user configuration. Ports that are configured as auto-VFL ports will run the discovery protocol which will automaticlly configure VFL IDs and member ports as well as chassis IDs.

Benefits of automatic Virtual Chassis:

- Existing switches configured in standalone mode will be unchanged and remain in standalone mode.
- Existing switches configured as part of a VC will be unchanged and remain as part of an existing VC.
- Newly shipped switches or switches with no configuration will default to automatic VC mode and the automatic VC feature will run.
- 45-day demo license enables the feature by default.
- By default the following ports are auto-VFL ports allowing newly installed switches to participate in the automatic VC feature:
    - OS10K - The first port of each 10G or 40G module
    - OS6900-X and T models - The last 5 ports of each chassis, including expansion slots. Ports without a transceiver present ARE included when determining VFL port eligibility.

- OS6900-Q32 - The last 5 ports among ports 28-32. Ports with a 40G-to-10G splitter cable inserted are counted as four ports. (i.e. 32a, 32b, 32c, 32d)

Note: – The OS6900-Q32 will by default have a higher chassis priority (120) than other OS6900 models. When creating a new mixed VC using default settings the OS6900-Q32 will become the Master chassis when combined with other OS6900 models.

Automatic Remote Configuration

The Automatic Remote Configuration capability automates and simplifies the deployment of network installations eliminating the need for manual configuration of each switch. It also ensures that each switch is compliant with the centrally controlled device configuration policies and firmware revisions.

This feature allows a newly deployed OmniSwitch to automate the process through an instruction file that provides the necessary actions to download its configuration and any necessary firmware upgrades with no user intervention by doing the following:

1. Automatically configures the switch with an IP client interface through various methods.
2. Lease an IP address, mask, default gateway, and system name from a reachable DHCP server.
3. Download an instruction file with information to obtain the configuration file, image files and/or script files from a given TFTP, FTP or SCP servers.
4. Download and apply the image and configuration file.
5. Automatically reboot with the upgraded image files and switch configuration file or if no images or boot configuration is downloaded scripted instructions are executed on the fly and the switch is made available remotely.

Automatic Fabric

The Automatic Fabric feature reduces the burden of configuration on the administrator. Dynamic recognition of the neighboring elements will allow for quick, out-of-the box configuration and reduced administrative overhead. Automatic fabric can be used to dynamically discover and configure a switch for the LACP, SPB, MVRP, and IP protocols and is supported in both standalone or virtual chassis mode. Once the automatic fabric feature begins the following occurs:

1. The switch will first attempt LACP discovery for a set discovery window.
2. After the LACP discovery window expires, SPB auto discovery will occur.
3. After SPB discovery window expires, MVRP auto discovery will occur.
4. The automatic IP discovery runs at the same time with the other protocols.

SPB Automatic Fabric Enhancements

Automatic fabric supports loopback detection which allows for the detection and breaking of loops on SAP interfaces. Dual homed connections can be done through linkagg connections to two or more devices that are part of the same VC. If a switch is connected to multiple devices that are not part of the same VC, the port is converted to an access port and a loop detection protocol will be run on these ports.

During the SPB discovery all the ports are considered as network ports. The network port is converted to a UNP port and LBD is enabled only if at least one successful SPB adjacency is formed during the SPB discovery.

Service profiles can be configured and applied for SAP ports or interface, profiles used for SAP:

- Single service profile: Defines single default service SAP profile. Only untagged traffic on the UNP-SPB port is learned in this profile.
- Auto-VLAN service profile: The incoming traffic is automatically learned. The SAP bindings for the concerned VLANs are automatically created based on the traffic sensed.

The SAP profile can be configured globally or set for a specific port or range of ports for the switch.

IP Protocol Auto Configuration
This enhancement extends automatic fabric feature to IP protocols. The IP protocols supported are OSPFv2, OSPFv3, ISIS IPv4 and ISIS IPv6. The automatic configuration of IP Protocols and automatic configuration of other protocols supported by automatic fabric can run in parallel.

Automatic Configuration of IP Protocols reduces the burden of configuration of neighboring routers on the administrator. Dynamic recognition of the neighboring routers will allow for quick out-of-the box configurations.

Automatic IP runs only when an active IP interface exists on the switch and automatic IP is enabled. Once an IP interface is created the interface will listen for hello packets from the neighboring devices and automatically configure the basic routing parameters based on the information received in the hello packets.

Saving the discovered configuration

Manually - The configuration discovered will be saved to the configuration file after discovery if the **write memory** command is given.

Automatically - The system will save the discovered configuration to the configuration file at set periods. This feature can be enabled or disabled and the interval changed as desired.

Management Feature Descriptions

Embedded Python Scripting
The OmniSwitch includes many standard Python packages to access AOS and system functions. This feature allows administrators to create Python scripts and associate these scripts with specific traps. When the traps are generated by the switch, the pre-configured scripts will be run on the switch. This provides the capability to adapt to a dynamically changing network and customize how the switch should react to those changes. There are multiple ways to execute Python on the switch:

- Automatically, as an event-action when a trap occurs
- Interactively, from the console
- In a script file executed by command from the console

AOS Python includes many standard Python packages for:

- OS access and issuing AOS commands
- Sending email and database access

IP Managed Services
By default, most applications that run on IP use the egress IP interface address as the source IP, while using a socket to communicate with a peer/server. However, it may be desirable to have some applications use a specific source IP for the packets that are sent out using the socket. This feature provides ('ip service source-ip' command) the ability to configure a permanent source IP interface to send packets. The source IP interface can be the Loopback0 address or user defined IP interface.

OpenFlow

OpenFlow is a communications interface defined between the control and forwarding layers that is used in a Software Defined Network (SDN). OpenFlow separates the control plane and the data plane in the switch. Traditionally, switches and routers have made decisions on where packets should travel based on rules local to the device. With OpenFlow, only the data plane exists on the switch itself, and all control decisions are communicated to the switch from a central Controller. If the device receives a packet for which it has no flow information, it sends the packet to the Controller for inspection, and the Controller determines where that packet should be sent based on QoS-type rules configured by the user (drop the packets to create a firewall, pass the packets to a specific port to perform load balancing, prioritize packets, etc).

The OmniSwitch can operate in AOS or OpenFlow mode, including a modified OpenFlow mode known as Hybrid mode. AOS will designate the ports managed/controlled by AOS or by OpenFlow on a per-port basis. By default, ports are managed/controlled by AOS.

OpenFlow 1.0 and 1.3.1 are supported. The following are the key components available for OpenFlow support.

- OpenFlow Logical Switch - An OpenFlow logical switch consists of a portion of the switch's resources that are managed by an OpenFlow Controller (or set of Controllers) via the OpenFlow Agent. Up to 3 logical switches can be configured with each switch supporting up to three controllers. A logical switch has a VLAN, physical ports, and/or link aggregate ports assigned to it. All packets received on these ports are forwarded directly to the OpenFlow agent. Spanning tree and source learning do not operate on OpenFlow assigned ports.
- OpenFlow Normal Mode - In Normal Mode, the logical switch operates as per the OpenFlow standards.
- OpenFlow Hybrid Mode (API) - In Hybrid mode, the logical switch acts as an interface through which the Controller may insert flows. These flows are treated as QoS policy entries and offer the same functionality. A Hybrid mode logical switch operates on all ports, link aggregates, and VLANs not assigned to other OpenFlow logical switches. Only one logical switch can be active in Hybrid mode.

Note: OpenFlow is supported on both standalone and Virtual Chassis.

Security Feature Descriptions

802.1x

Physical devices attached to a LAN port on the switch through a point-to-point LAN connection can be authenticated through the OmniSwitch using port-based network access control. This control is available through the IEEE 802.1X standard implemented on the switch.

The Universal Network Profile (UNP) functionality uses this implementation of 802.1X to provide configurable device classification policies for authenticating both 802.1x clients (supplicants) and non-802.1x clients (non-supplicants). Such policies also determine the VLAN or service assignment of a device and are particularly useful for providing secure network access to guest clients.

The 802.1x authentication process uses the Extensible Authentication Protocol (EAP) between an end device and a network device to authenticate the supplicant (802.1x device) through a RADIUS server. If authentication returns a UNP, the supplicant is assigned to that UNP. If a UNP name is not returned or authentication fails, then the UNP port and classification rule configuration provides the network access control for the supplicant.

802.1x is supported for VLANs, SPBM, and VXLAN Services

SNMP Traps

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

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

No.	Trap Name	Description
		link because of excessive errors.
19	portViolationTrap	This trap is sent when a port violation occurs. The port violation trap will indicate the source of the violation and the reason for the violation.
20	dvmrpNeighborLoss	A 2-way adjacency relationship with a neighbor has been lost. This trap is generated when the neighbor state changes from "active" to "one-way," "ignoring" or "down." The trap is sent only when the switch has no other neighbors on the same interface with a lower IP address than itself.
21	dvmrpNeighborNotPruning	A non-pruning neighbor has been detected in an implementation-dependent manner. This trap is generated at most once per generation ID of the neighbor. For example, it should be generated at the time a neighbor is first heard from if the prune bit is not set. It should also be generated if the local system has the ability to tell that a neighbor which sets the prune bit is not pruning any branches over an extended period of time. The trap should be generated if the router has no other neighbors on the same interface with a lower IP address than itself.
22	risingAlarm	An Ethernet statistical variable has exceeded its rising threshold. The variable's rising threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
23	fallingAlarm	An Ethernet statistical variable has dipped below its falling threshold. The variable's falling threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
24	stpNewRoot	Sent by a bridge that became the new root of the spanning tree.
25	stpRootPortChange	A root port has changed for a spanning tree bridge. The root port is the port that offers the lowest cost path from this bridge to the root bridge.
26	mirrorConfigError	This trap is sent when any NI fails to configure mirroring. Due to this error, port mirroring session will be terminated.
27	mirrorUnlikeNi	The mirroring configuration is deleted due to the swapping of different NI board type. The Port Mirroring session which was active on a slot cannot continue with the insertion of different NI type in the same slot.
28	slbTrapOperStatus	A change occurred in the operational status of the server load balancing entity.
29	sessionAuthenticationTrap	An authentication failure trap is sent each time a user

No.	Trap Name	Description
		authentication is refused.
30	trapAbsorptionTrap	The absorption trap is sent when a trap has been absorbed at least once.
31	alaDoSTrap	Indicates that the sending agent has received a Denial of Service (DoS) attack.
32	ospfNbrStateChange	Indicates a state change of the neighbor relationship.
33	ospfVirtNbrStateChange	Indicates a state change of the virtual neighbor relationship.
34	lnkaggAggUp	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.
35	lnkaggAggDown	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
36	lnkaggPortJoin	This trap is sent when any given port of the link aggregate group goes to the attached state.
37	lnkaggPortLeave	This trap is sent when any given port detaches from the link aggregate group.
38	lnkaggPortRemove	This trap is sent when any given port of the link aggregate group is removed due to an invalid configura tion.
39	monitorFileWritten	This trap is sent when the amount of data requested has been written by the port monitoring instance.
40	alaVrrp3TrapProtoError	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
41	alaVrrp3TrapNewMaster	The SNMP agent has transferred from the backup state to the master state.
42	chassisTrapsPossibleDuplicateMac	The old PRIMARY element cannot be detected in the stack. There is a possiblity of a duplicate MAC address in the network
43	lldpRemTablesChange	A lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes.
44	pimNeighborLoss	A pimNeighborLoss notification signifies the loss of an adjacency with a neighbor.
45	pimInvalidRegister	An pimInvalidRegister notification signifies that an invalid PIM Register message was received by this device
46	pimInvalidJoinPrune	A pimInvalidJoinPrune notification signifies that an invalid PIM Join/Prune message was received by this device.

No.	Trap Name	Description
47	pimRPMappingChange	An pimRPMappingChange notification signifies a change to the active RP mapping on this device.
48	pimInterfaceElection	An pimInterfaceElection notification signifies that a new DR or DR has been elected on a network.
49	pimBsrElectedBSRLostElection	This trap is sent when the current E-BSR loses an election to a new Candidate-BSR.
50	pimBsrCandidateBSRWinElection	This trap is sent when a C-BSR wins a BSR Election.
51	lpsViolationTrap	A Learned Port Security (LPS) violation has occurred.
52	lpsPortUpAfterLearningWindowExpiredT	When an LPS port joins or is enabled after the Learning Window is expired, the MAC address learning on the port will be disabled, and this trap is generated as a notification.
53	lpsLearnMac	Generated when an LPS port learns a bridged MAC.
54	gvrpVlanLimitReachedEvent	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
55	alaNetSecPortTrapAnomaly	Trap for an anomaly detected on a port.
56	alaNetSecPortTrapQuarantine	Trap for an anomalous port quarantine.
57	ifMauJabberTrap	This trap is sent whenever a managed interface MAU enters the jabber state.
58	udldStateChange	Generated when the state of the UDLD protocol changes.
59	ndpMaxLimitReached	This IPv6 Trap is sent when the hardware table has reached the maximum number of entries supported. Not supported in AOS Release 7.
60	ripRouteMaxLimitReached	This trap is sent when the RIP database reaches the supported maximum number of entries. When the maximum number is reached, RIP discards any new updates.
61	ripngRouteMaxLimitReached	This trap is sent when the RIPng database reaches the supported maximum number of entries. When the maximum number is reached, RIPng discards any new updates.
62	alaErpRingStateChanged	This trap is sent when the ERP ring State has changed from “Idle” to “Protection”.
63	alaErpRingMultipleRpl	This trap is sent when multiple RPLs are detected in the ring.
64	alaErpRingRemoved	This trap is sent when the ring is removed dynamically.
65	ntpMaxAssociation	This trap is generated when the maximum number of

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

No.	Trap Name	Description
79	multiChassisStpBlockingStatus	This trap is sent to indicate the STP status for some VLANs on the VFLink is in blocking state.
80	multiChassisLoopDetected	This trap is sent to indicate a loop has been detected.
81	multiChassisHelloTimeout	This trap is sent to indicate the hello timeout has occurred.
82	multiChassisVflinkDown	This trap is sent to indicate the VFLink is down.
83	multiChassisVFLMemberJoinFailure	This trap is sent to indicate a port configured as virtual-fabric member is unable to join the virtual-fabric link.
84	alaDHLVlanMoveTrap	When linkA or linkB goes down or comes up and both ports are are part of some vlan-map, this trap is sent to the Management Entity, with the DHL port information.
85	alaDhcpClientAddressAddTrap	This trap is sent when a new IP address is assigned to DHCP Client interface.
86	alaDhcpClientAddressExpiryTrap	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address.
87	alaDhcpClientAddressModifyTrap	This trap is sent when the DHCP client is unable to obtain the existing IP address and a new IP address is assigned to the DHCP client
88	vRtrIsisDatabaseOverload	This notification is generated when the system enters or leaves the overload state.
89	vRtrIsisManualAddressDrops	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
90	vRtrIsisCorruptedLSPDetected	This notification is generated when an LSP that was stored in memory has become corrupted.
91	vRtrIsisMaxSeqExceedAttempt	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
92	vRtrIsisIDLenMismatch	A notification sent when a PDU is received with a different value of the System ID Length.
93	vRtrIsisMaxAreaAddrsMismatch	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
94	vRtrIsisOwnLSPPurge	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
95	vRtrIsisSequenceNumberSkip	When an LSP is received without a System ID and different contents.
96	vRtrIsisAutTypeFail	A notification sent when a PDU is received with the wrong authentication type field.

No.	Trap Name	Description
97	vRtrIsisAuthFail	A notification sent when a PDU is received with an incorrent authentication information field.
98	vRtrIsisVersionSkew	A notification sent when a a Hello PDU is received from an IS running a different version of the protocol.
99	vRtrIsisAreaMismatch	A notification sent when a Hello PDU is received from an IS which does not share any area address.
100	vRtrIsisRejectedAdjacency	A notification sent when a Hello PDU is received from an IS, but does not establish an adjacency due to a lack of resources.
101	vRtrIsisLSPTooLargeToPropagate	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
102	vRtrIsisOrigLSPBufSizeMismatch	A notification sent when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for the originating L1LSP BufferSize or originating L2LSPBufferSize respectively. Also when a Level 1 LSP or Level2 LSP is received containing the originating LSPBufferSize option and the value in the PDU option field does not match the local value for originating L1LSP BufferSize or originatingL2LSP BufferSize respectively.
103	vRtrIsisProtoSuppMismatch	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
104	vRtrIsisAdjacencyChange	A notification sent when an adjacency changes state, entering or leaving state up. The first 6 bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
105	vRtrIsisCircIdExhausted	A notification sent when ISIS cannot be started on a LAN interface because a unique circId could not be assigned due to the exhaustion of the circId space.
106	vRtrIsisAdjRestartStatusChange	A notification sent when an adjancency's graceful restart status changes.
107	mvrpVlanLimitReachedEvent	This trap is sent when the number of VLANs learned dynamically by MVRP reaches the configured limit.
108	alaHAVlanClusterPeerMismatch	The trap is sent when parameteras configured for this cluster ID (Level 1 check) does not match accross the MCLAG peers.
109	alaHAVlanMCPeerMismatch	The trap is sent when when the cluster parameters are matching on the peers, but MCLAG is not configured or clusters are not in operational state.
110	alaHAVlanDynamicMAC	The trap is sent when the dynamic MAC is learned on non-server cluster port
111	unpMcLagMacIgnored	This trap is sent when a MAC/User is dropped because the VLAN does not exist or UNP is not enabled on the

No.	Trap Name	Description
		MCLAG.
112	unpMcLagConfigInconsistency	This trap is sent when a configuration becomes “Out of Sync".
113	multiChassisGroupConsisFailure	This trap is sent when there is an inconsistency between local and peer chassis group.
114	multiChassisTypeConsisFailure	This trap is sent when there is an inconsistency between local and peer chassis group.
115	alaPimNonBidirHello	This trap is sent when a bidir-capable router has received a PIM hello from a non-bidir-capable router. It is generated whenever the counter alaPimsmNon-BidirHelloMsgsRcvd is incremented, subject to the rate limit specified by alaPimsmNonBidirHelloNotificationPeriod.
116	dot1agCfmFaultAlarm	This trap is sent when a MEP has a persistent defect condition. A notification (fault alarm) is sent to the management entity with the OID of the MEP that has detected the fault.
117	alaSaaIPIterationCompleteTrap	This trap is sent when an IP SAA iteration is completed.
118	alaSaaEthIterationCompleteTrap	This trap is sent when when an eth-LB or Eth-DMM SAA iteration is completed.
119	alaSaaMacIterationCompleteTrap	This trap is sent when a MAC iteration is complete.
120	virtualChassisStatusChange	This trap is sent when a chassis status change is detected.
121	virtualChassisRoleChange	This trap is sent when a chassis role change is detected.
122	virtualChassisVflStatusChange	This trap is sent when a vflink status change is detected.
123	virtualChassisVflMemberPortStatusCh	This trap is sent when a vflink member port has a change of status.
124	virtualChassisVflMemberPortJoinFail	This trap is sent when a port configured as virtual-fabric member is unable to join the virtual-fabric link.
125	lldpRemTablesChange	This trap is sent when the value of lldpStatsRemTablelastChange Time changes. It can be utilized by an NMS to trigger LLDP remote systems table maintenance polls.

No.	Trap Name	Description
126	vRtrLdpInstanceStateChange	This trap is sent when the LDP module changes state either administratively or operationally.
127	evbFailedCdcpTlvTrap	This trap is sent when bridge receives a CDCP packet with: - Wrong TLV type, or - Wrong OUI, or - Role is set to Bridge, or - Wrong default channel(scid), or - Incorrect channel number(scid).
128	evbFailedEvbTlvTrap	This trap is sent when bridge receives an EVBTLV packet with: - Wrong TLV type. or - Incorrect TLV length, or - Wrong OUI.
129	evbUnknownVsiManagerTrap	This trap is sent when bridge receives a VDP packet with: - Unknown Manager ID type, or - Wrong Manager ID length.
130	evbVdpAssocTlvTrap	This trap is sent when bridge receives an ASSOC TLV in a VDP packet with: - Null VID found and number of entry field is not 1, or - Unknown filter format, or - Null VID on De-Assoc TLV type, or - VSI included more than Max number of filter info entries
131	evbCdcpLldpExpiredTrap	This trap is sent when an LLDP Timer expired in the bridge. The timer expires when LLDP does not receive CDCP TLV within a specified interval.
132	evbTlvExpiredTrap	This trap is sent when an LLDP Timer expired in the bridge. The timer expires when LLDP doesn't not receive EVB TLV within a specified interval.
133	evbVdpKeepaliveExpiredTrap	This trap is sent when a VDP Keep Alive Timer expired in bridge. The timer expires when the bridge doesn't not receive VDP Keep Alive message within a specified interval.
134	smgrServiceError	This trap is sent when there is a failure to create/delete a service.
135	smgrServiceHwError	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a

No.	Trap Name	Description
		service, or to program the hardware tables for a service.
136	smgrServiceSapError	This trap is sent when there is a failure to create/delete a Service Access Point.
137	smgrServiceSapHwError	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a SAP, or to program the hardware tables for a SAP.
138	smgrServiceSdpError	This trap is sent when there is a failure to create/delete a Service Distribution Point.
139	smgrServiceSdpHwError	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP, or to program the hardware tables for an SDP.
140	smgrServiceSdpBindError	This trap is sent when there is a failure to create/delete an SDP Bind.
141	smgrServiceSdpBindHwError	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP Bind, or to program the hardware tables for an SDP Bind.
142	smgrGeneralError	This trap is sent when there is a general system failure detected during normal system operation.
143	smgrStatusChange	This trap is sent when there is a status change for a group of selected services.
144	portViolationNotificationTrap	This trap is sent when a port violation is cleared.
145	multiChassisConsisFailureRecovered	This trap is sent when the system has recovered from a multi-chassis inconsistency between the local and peer switches
146	alaSaaPacketLossTrap	This trap is sent when a packet is lost during a test.
147	alaSaaJitterThresholdYellowTrap	This trap is sent when the Jitter Threshold crosses 90%.
148	alaSaaRTTThresholdYellowTrap	This trap is sent when the RTT Threshold crosses 90%.
149	alaSaaJitterThresholdRedTrap	This trap is sent when the Jitter threshold is crossed.
150	alaSaaRTTThresholdRedTrap	This trap is sent when the RTT threshold is crossed.
151	chassisTrapsDuplicateMacClear	This trap is sent when the old Master Chassis has rejoined the Virtual Chassis as a slave.
152	alaFipsConfigFilterResourceLimit	The allowed maximum percentage of filter resources configured from the allocated FIPS resources is exceeded.
153	virtualChassisUpgradeComplete	Critical trap indicates whether the software upgrade process has failed after a timeout or completed successfully. Note that if the process fails, it may be still possible for the system to recover if the process successfully completes later after the expired timeout.
154	appFPSignatureMatchTrap	This trap is sent when a traffic flow matches an application signature.

No.	Trap Name	Description
155	virtualChassisVflSpeedTypeChange	This trap indicates whenever vfl speed is changed.
156	alaSIPSnoopingACLPreemptedBySOSCall	This trap is sent when a SIP snooping RTP/RTCP ACL entry is preempted by an SOS call.
157	alaSIPSnoopingRTCPOverThreshold	This trap is sent when one or more RTCP parameters exceeds the threshold limit.
158	alaSIPSnoopingRTCPPktsLost	This trap is sent when RTCP packets are lost due to rate limiting.
159	alaSIPSnoopingSignallingLost	This trap is sent when when SIP signalling messages are lost due to rate limiting.
160	alaSIPSnoopingCallRecordsFileMoved	This trap is generated when SIP SNOOPING ended call records flash file is moved from /flash/switch/sip_call_record.txt to /flash/switch/sip_call_record.txt.old. This happens when the configured call record storage limit is reached and possibly at boot-up if /flash/switch/sip_call_record.txt from previous run exists at the first check.
161	alaIPv6NeighborLimitExceeded	This trap is sent when the system-wide neighbor cache limit is exceeded.
162	alaIPv6NeighborVRFLimitExceeded	This trap is sent when a per-VRF neighbor cache limit is exceeded.
163	alaIPv6InterfaceNeighborLimitExceed	This trap is sent when a per-interface neighbor cache limit is exceeded.
164	alaDyingGaspTrap	This trap is sent when a switch has lost all power.
165	alaDhcpSrvLeaseUtilizationThreshold	This trap is sent when the lease utilization on a subnet exceeds or falls below the configured threshold value.
166	alaDHCPv6SrvLeaseUtilizationThresho	This trap is sent when the lease utilization on a subnet exceeds or falls below the configured threshold value.
167	smgrServiceStatusChange	This trap is sent when there is a change in service operating status. A service is operationally up when it's admin-up and there's at least one active SAP or one active bind that is operationally up.
168	smgrSapStatusChange	This trap is sent when there is a change in SAP operating status. A SAP is operationally up when it's admin-up and the link status of the physical or logical port of the SAP is operationally up.
169	smgrSdpStatusChange	This trap is sent when there is a change in SDP operating status. For SPB, the SDP is dynamically created or destroyed as calculated by ISIS protocol when a unicast/multicast path to reach a neighbor node is determined.

No.	Trap Name	Description
170	smgrSdpBindStatusChange	This trap is sent when there is a change in SDP Bind operating status. For SPB, the SDP Bind is dynamically created or destroyed as detected by ISIS when the same ISID is configured in the neighbor node.
171	alaPethPwrSupplyConflictTrap	Power supply type conflict trap.
172	pethPwrSupplyNotSupportedTrap	Power supply not supported trap.
173	chasTrapsBPSLessAllocSysPwr	This trap is sent when there is insufficient system power being provided by the BPS.
174	chasTrapsBPSStateChange	This trap is sent when a BPS power supply is inserted or removed.
175	chasTrapsNiBPSFETStateChange	This trap is sent when there is a BPS FET state change.
176	alaDhcpBindingDuplicateEntry	This trap is sent when when ther is MAC Movement in DHCP-Binding Table.
177	alaVCSPProtectionTrap	This trap is sent when a virtual chassis enters the split protection state.
178	alaVCSPRecoveryTrap	This trap is sent when a split virtual chassis enters the active state.
179	pethPsePortOnOffNotification	Indicates if power inline port is or is not delivering power to the a power inline device.
180	pethMainPowerUsageOnNotification	Indicates that the power inline usage is above the threshold.
181	pethMainPowerUsageOffNotification	Indicates that the power inline usage is below the threshold.
182	chasTrapsBPSFwUpgradeAlert	This trap is sent when a BPS firmware upgrade is required.
183	alaAppMonAppRecordFileCreated	This trap is sent after the application records monitored in the past hour are written to the flash file.
184	alaAppMonFlowRecordFileCreated	This trap is sent after the pre-configured number of application monitoring flow records are written to the flash file.
185	alaDPIFlowRecordFileCreated	This trap is sent after the pre-configured number of deep packet inspection flow records are written to the flash file.
186	alaLbdStateChangeToShutdown	This trap is sent when a port is shut down.
187	alaLbdStateChangeForClearViolationA	This trap is sent when the port state changes from shutdown due to “clear-violation-all”.
188	alaLbdStateChangeForAutoRecovery	This trap is sent when a port state changes from shutdown due to the auto-recovery mechanism.

No.	Trap Name	Description
189	alaAutoConfigAutoFabricEnableTrap	This object specifies the threshold status of subnet utilization (Yes/No).
190	alaVMSnoopingVMLearntAlert	This trap is sent when a new Virtual Machine is learned by the system.
191	alaVMSnoopingVMRemovedAlert	This trap is sent when virtual machine entry ages out or is removed from the database.
192	alaVMSnoopingReservedHwResourceLimi	This trap is sent when the reserved hardware resource reaches a cutoff limit.
193	alaDistArpItfChange	This trap is sent when an interface is re-assigned to a new designated NI.
194	alaDistArpNiThreshold	This trap is sent when the number of ARPs in hardware has reached the reassignment threshold.
195	smgrVxlanSdpBindStatusChange	This trap is sent when there is an change in SDP Bind operating status. An SDP Bind is dynamically created when a VTEP is discovered, or destroyed when the reachability to the VTEP is gone.

Unsupported Software Features

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

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

Unsupported CLI Commands

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

Software Feature	Unsupported CLI Commands
Chassis	reload slot
SLB	server-cluster port all

Open Problem Reports and Feature Exceptions

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

DCB

PR	Description	Workaround
203454	On an OS6900-Q32 a custom profile can only be created by importing the pre-configured DCB-8 profile.	There is no known workaround at this time.

Hardware

PR	Description	Workaround
198986	The "show interfaces" command does not display the link-quality for any of the lanes on an OS6900-Q32 40G port.	There is no known workaround at this time.

Layer 3

PR	Description	Workaround
204426	If a tagged router port is first created for any physical port, then no other router port is allowed to be configured on the same physical port.	Always create an untagged router port first if any tagged router ports need to be created on the same physical port.

Multicast

PR	Description	Workaround
193941	The IPMS source timeout (30 sec default) can be manipulated by an IGMP client sending a combination of join and leave messages right before the timer expires causing the multicast source to continue to live on and never expire.	There is no known workaround at this time.

SPB

PR	Description	Workaround
196295	IP-SPB traffic is not being forwarded across the VFL on an OS10K when the VFL is configured on an OS10K-XNI-U32S.	It is not recommended to have the VFL configured on an OS10K-XNI-U32S with SPB configured. Configure the VFL on an OS10K-XNI-U32E instead.
202168	After flushing all dynamic MACs using the 'mac-learning flush dynamic' command, any newly learned MACs are never displayed again on an SPB domain.	Use the 'mac-learning flush domain-spb dynamic' to flush the SPB domain.

UNP

PR	Description	Workaround
203461	It can take up to 90 seconds in some cases for UNP users to be relearned in case of an edge port toggle.	There is no known workaround at this time.

Webview

PR	Description	Workaround
196794	In Webview, no local queriers are detected under "IPv4 Multicast Detected Queriers".	Use the 'show ip multicast querier' command to view the current queriers.

Virtual Chassis

PR	Description	Workaround
198541	After a reload or CMM takeover on an OS10K, the EMP (Chassis and VC) is down as seen from the 'show ip interface' command on console. Ping and telnet to the EMP address fail as a result.	In the event EMP is down but console access is still available, from su, type: 'ip link set eth3 down' followed by 'ip link set eth3 up' This will restore EMP interface status.
204129	Issue is seen when EMP-VC IP address is configured without a chassis EMP IP address. When EMP cable is removed before a reboot and inserted after a reboot, the EMP-VC becomes unreachable though EMP-VC IP interface indicates it is UP. An iprm_0 intf4 error message is seen on the console : " iprmIntfEnable: Failed to find IPv4 interface 4106EMP"	Un-configure the EMP-VC interface and re-configure it again.
203912	When a factory default OS6900-Q32 with Demo Advanced license and an OS6900-X20/X40, that has both a permanent Advanced and a permanent Data Center license, are connected and booted together, the Virtual Chassis won't become operational due to a Mismatch-License-Config.	The chassis with permanent Data Center license should be configured such that it always becomes the Master.

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.

- 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
OS10K-GNI-C48E	OS10K-GNI-C48E
OS10K-GNI-U48E	OS10K-GNI-U48E
OS10K-XNI-U32S	OS10K-XNI-U32S
OS10K-XNI-U16L	OS10K-XNI-U16L
OS10K-XNI-U16E	OS10K-XNI-U16E
OS10K-XNI-U32E	OS10K-XNI-U32E
OS10K-QNI-U4E	OS10K-QNI-U4E
OS10K-QNI-U8E	OS10K-QNI-U8E

OS10K 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 replacement.
4. Insert replacement module of same type.
5. Wait for a message similar to the following to display on the console or issue the -> show module status command and wait for operational status to show 'UP':

   ChassisSupervisor niMgr info message:

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

Technical Support

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

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

Email: esd.support@alcatel-lucent.com

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

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

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

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

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

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

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: 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 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.

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.

Staggered Upgrade - A staggered upgrade is similar to ISSU but is designed for those situations that do not completely support ISSU. A staggered upgrade may be required when upgrading between different AOS release trees (i.e. 7.3.2 to 7.3.4) due to underlying code variations between the two releases which may not allow CMMs or Master/Slave chassis to communicate after one is upgraded to the newer version of code.

A staggered upgrade requires a script file to be run prior to the upgrade. The script will copy the required configuration and image files to the CMMs or chassis to be upgraded. It also provides a mechanism to allow the Primary CMM or Master chassis to know the upgrade has been completed successfully on the redundant CMM or Slave chassis before rebooting. This allows for an upgrade between different AOS release trees with minimal network disruption.

Supported Upgrade Paths and Procedures

	Upgrading From 7.3.3	Upgrading From 7.3.2	Upgrading From 7.3.1
OS6900 – VC	ISSU – Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported
OS6900 – Standalone	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported
OS10K – VC	N/A	ISSU – Not Supported Staggered Upgrade – Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported
OS10K – Standalone	N/A	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported	ISSU – Not Supported Staggered Upgrade – Not Supported Standard Upgrade - Supported

- If upgrading a standalone chassis or VC using a standard upgrade procedure please refer to Appendix B for specific steps to follow.
- If upgrading an OS6900 VC from 7.3.3.R01 to 7.3.4.R01 using ISSU please refer to Appendix C for specific steps to follow.
- If upgrading an OS10K or OS6900 VC using a staggered upgrade please refer to Appendix D for specific steps to follow to help minimize any network disruption.

Prerequisites

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

- Be the responsible party for maintaining the switch's configuration.
- Be aware of any issues that may arise from a network outage caused by improperly loading this code.
- Understand that the switch must be rebooted and network 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
```

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.

Appendix B: Standard Upgrade - OmniSwitch 6900/10K Standalone/Virtual Chassis

These instructions document how to upgrade an OS6900 or OS10K standalone or virtual chassis to 7.3.4.R01 using the standard upgrade procedure. Upgrading to 7.3.4.R01 using the standard upgrade procedure consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support website and download and unzip the 7.3.4.R01 upgrade files for the appropriate model. The archives contain the following:

- OS6900 Image Files **-** Tos.img
- OS10K Image Files – Ros.img, Reni.img

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.

2. 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.

3. 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** commmand**.**

```
OS6900-> show microcode
   /flash/working
   Package          Release                Size     Description
-----------------+------------------------+--------+-----------------------------------
Tos.img          7.3.4.450.R01            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.

4. 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 C: ISSU – OmniSwitch OS6900 Virtual Chassis

These instructions document how to upgrade an OS6900 virtual chassis to AOS release 7.3.4.R01 using ISSU. Upgrading a VC of OS6900 Switches to 7.3.4.R01 consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support Website and download and unzip the 7.3.4.R01 ISSU upgrade files for the OS6900. The archive contains the following:

- OS6900 Image Files **-** Tos.img
- ISSU Version File – issu_version

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
*** Add CLI output ***
```

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.

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** commmand**.**

```
OS6900-> show microcode
   /flash/working
   Package          Release                Size     Description
-----------------+-----------------------+--------+---------------------------------
Tos.img          7.3.4.450.R01            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 D: Staggered Upgrade - OmniSwitch OS10K/OS6900

These instructions document how to upgrade an OS10K or OS6900 VC to 7.3.4.R01 using a staggered upgrade process. Upgrading an OmniSwitch to 7.3.4.R01 using a staggered upgrade procedure consists of the following steps. The steps should be performed in order.

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support Website and download and unzip the 7.3.4.R01 upgrade files for the appropriate model. The archives contain the following:

- OS6900 Image Files – Tos.img
- OS10K Image Files **-** Ros.img, Reni.img
- Upgrade Script – vcof2-upgrade

2. Create a directory to hold the upgrade files on the Master chassis

```
OS10K-> mkdir /flash/issu_dir
```

3. FTP the upgrade files to the directories below on the Master chassis:

- Ros.img and Reni.img - /flash/issu_dir
- vcof2-upgrade - /flash

4. Execute the script on the Master chassis:

```
OS10K-> chmod a+x /flash/vcof2-upgrade
OS10K-> /flash/vcof2-upgrade  issu_dir
```

The above command sequence copies the **vcboot.cfg** and **vcsetup.cfg** from the current *Running* directory to **/flash/issu_dir** directory. It also copies the image files and config files to the secondary and Slave CMMs. It then creates the special upgrade helper file "**vcupgrade.cfg**" and copies it to the Slave. It then initiates a reload on the Slave with the new software which begins the upgrade process.

5. Verify the Software Upgrade

To verify that the software was successfully upgraded to 7.3.4.R01, use the **show microcode** command as shown below:

```
OS10K-> show microcode
/flash/working

Package        Release           Size       Description
---------------+-----------------+----------+-----------------
Ros.img        7.3.4.450.R01     106031376  Alcatel-Lucent OS
Reni.img       7.3.4.450.R01     106031376  Alcatel-Lucent OS
```

6. 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: Previous Release Feature Summary

Existing Hardware/Software Feature Summary – AOS 7.3.3

Feature	Platform	License
Hardware Support		
- OS-XNI-U12E	OS6900	Base
- SFP-FC-SR Transceiver	OS6900	Base
Data Center Feature Support		
- FCoE/FC Gateway	6900	Data Center
- CEE DCBX Version 1.01	6900	Data Center
Layer 3 Feature Support		
- ISIS – IPv4/IPv6	6900	Advanced
- BGP 4-Octet ASN	6900	Advanced
Management		
- Virtual Chassis mesh of 6 chassis with ISSU support	6900	Advanced
Early Availability Feature Support		
- OpenFlow Agent versions 1.3.1 and 1.0 (Normal and Hybrid modes)	6900	Base
- Internal IPv4/IPv6 DHCP Server	6900	Base
- OmniSwitch Networking Plug-in for OpenStack	6900	Base
- M-ISIS	6900	Advanced

Existing Hardware/Software Feature Summary – AOS 7.3.2.R01

Feature	Platform	License
Hardware Feature Support		
- OmniSwitch 6900-T20		
- OmniSwitch 6900-T40		
- OS-XNI-T8		
Data Center Feature Support		
- FIP Snooping	OS10K/6900	Data Center
- Virtual Maching Performance Monitoring	OS10K/6900	Data Center
Layer 2 Feature Support		
- Dynamic Auto Fabric	OS10K/6900	Base
Layer 3 Feature Support		
- IPv4 over SPBM	OS10K/6900	Advanced
- Interop between PIM & DVMRP	OS10K/6900	Base
- Non-Contiguous Mask and IPv6 Gateway Support	OS10K/6900	Base
- Increase VRF Instances	OS10K/6900	Base
Management/Additional Feature Support		
- Command Abbreviation	OS10K/6900	Base
- Web Services & CLI Scripting	OS10K/6900	Base
- Enhanced Server & Session Limits	OS10K/6900	Base
Additional Feature Support		
- Application Fingerprinting	OS10K/6900	Base
- Fault Propagation and Link Flapping - Wait to Shutdown	OS10K/6900	Base

Existing Hardware/Software Feature Summary – AOS 7.3.1.R01

Feature	Platform	License
Hardware Feature Support		
OS10K-XNI-U16L OS10K-XNI-U16E OS10K-XNI-U32E OS10K-QNI-U4E OS10K-QNI-U8E QSFP-40G-LR Transceiver SFP-10G-24DWDM80 Transceiver SFP-10G-GIG-SR Transceiver		
Data Center Feature Support		
Shortest Path Bridging (SPB)	OS10K/6900	Advanced
Data Center Bridging • DCBX • ETS • PFC	 OS10K/6900 OS10K/6900 OS10K/6900	 Data Center Data Center Data Center
Edge Virtual Briding (EVB)	OS10K/6900	Data Center
Virtual Network Profiles • SAP/SPB-M Services • Customer Domains (Multi-tenancy) • Dynamic SAP • UNP over MC-LAG on OS10K	 OS10K/6900 OS10K/6900 OS10K/6900 OS10K/6900	 Base Base Base Base
Layer 2 Feature Support		
Ethernet Ring Protection v2 (ERPv2)	OS10K/6900	Base
Layer 3 Feature Support		
VRF Management	OS10K/6900	Base
VRF Route Leak	OS10K/6900	Base
Management Feature Support		
Virtual Chassis	OS10K/6900	Advanced

Feature	Platform	License
SFP+ Line Diags & Enhanced Port Performance (EPP)	OS10K/6900	Base
License Management	OS10K/6900	Base
Ethernet OAM • ITU Y1731 and 802.1ag	OS10K/6900 OS10K/6900	Base
Service Assurance Agent	OS10K/6900	Base

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

Existing Hardware/Software Feature Summary – AOS 7.2.1.R02

Feature	Platform	License
Hardware Feature Support		
OmniSwitch 6900 Rear-to-Front Cooling OS-QNI-U3 Module OS-HNI-U6 Module QSFP-40G-SR Transceiver QSFP-40G-C Transceiver OS6900-BP-R (YM-2451F) Power Supply OS6900-BPD-R (YM-2451P) Power Supply OS6900-FT-R FanTray		
Layer 2 Feature Support		
High Availability VLAN • Added support for OS10K • HA-VLAN with MCLAG	 OS10K OS10K/6900	 Base Base
Multi-Chassis Link Aggregation • Configurable Chassis Group ID (Multiple MC-LAG Domains) • Standalone Port in VIP VLAN • SLB Over MC-LAG	 OS10K/6900 OS10K/6900 OS10K/6900	 Base Base Base
MVRP • Added support for OS10K	 OS10K	 Base
Universal Network Profiles • UNP with Dynamic Profiles • UNP with Link-Aggregation • UNP with MC-LAG • UNP with Learned Port Security	 OS6900 OS6900 OS6900 OS6900	 Base Base Base Base
Layer 3 Feature Support		
16 ECMP routes for IPv6	OS10K/6900	Base
Qos		
VFC/VoQ Profiles • Added support for profiles 2-4 • Added support for WRED	 OS10K/6900 OS6900	 Base Base

Feature	Platform	License
Security		
Learned Port Security Enhancements	OS10K/6900	Base

Existing Hardware/Software Feature Summary – AOS 7.2.1.R01

Feature	Platform	License
Hardware Feature Support		
OmniSwitch 6900-X20 OmniSwitch 6900-X40 OS-XNI-U4 OS-XNI-U12 OS6900-BP-F (YM-2451C) Power Supply OS6900-BPD-F (YM-2451D) Power Supply OS6900-FT-F FanTray		
Manageability Feature Support		
CLI	OS6900	Base
Ethernet Interfaces	OS6900	Base
License Management	OS6900	Base
Multiple VRF Routing and Forwarding	OS6900	Advanced
Network Time Protocol (NTP)	OS6900	Base
Pause Control(RX) /Flow Control	OS6900	Base
Remote Access FTP SCP SSH/SFTP Telnet TFTP	OS6900	Base
Resiliency Features Hot Swap Expansion Modules Power Supply Redundancy Fan Redundancy	OS6900	Base
SNMP	OS6900	Base
Software Rollback – Multi-Image/Multi-Config	OS6900	Base
Storm Control	OS6900	Base
Text File Configuration	OS6900	Base
UDLD	OS6900	Base
USB Support	OS6900	Base
Web-Based Management (WebView)	OS6900	Base

Feature	Platform	License
Layer 2 Feature Support		
802.1AB with MED Extensions	OS6900	Base
802.1Q	OS6900	Base
Configurable Hash Mode	OS6900	Base
HA-VLAN	OS6900	Base
Link Aggregation –Static and LACP (802.3ad)	OS6900	Base
Multi-Chassis Link Aggregation	OS6900	Base
MVRP	OS6900	Base
Source Learning	OS6900	Base
Spanning Tree • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard	OS6900	Base
Universal Network Profiles (UNP)	OS6900	Base
VLANs	OS6900	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS6900	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS6900	Base
BGP4 with Graceful Restart	OS6900	Advanced
DNS Client	OS6900	Base
GRE	OS6900	Base
IP Multicast Routing	OS6900	Advanced
IP Multicast Switching (IGMP)	OS6900	Base
IP Multicast Switching (Proxying)	OS6900	Base
IP Multinetting	OS6900	Base
IP Route Map Redistribution	OS6900	Base

Feature	Platform	License
IP-IP Tunneling	OS6900	Base
OSPFv2	OS6900	Advanced
RIPv1/v2	OS6900	Base
Routing Protocol Preference	OS6900	Base
Server Load Balancing	OS6900	Base
VRRPv2	OS6900	Advanced
IPv6 Feature Support		
BGP4 BGP IPv6 Extensions	OS6900	Advanced
IPSec IPv6 OSPFv3 RIPng	OS6900	Advanced
IPv6 Client and/or Server Support	OS6900	Base
IPv6 Multicast Routing	OS6900	Advanced
IPv6 Multicast Switching (MLD v1/v2)	OS6900	Base
IPv6 Routing	OS6900	Advanced
IPv6 Scoped Multicast Addresses	OS6900	Base
IPv6 Neighbor Discovery Support	OS6900	Base
OSPFv3	OS6900	Advanced
RIPng	OS6900	Advanced
VRRPv3	OS6900	Advanced
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS6900	Base
Ingress and egress bandwith shaping	OS6900	Base
Policy Based Routing	OS6900	Advanced
Tri-Color Marking	OS6900	Base
Multicast Feature Support		
DVMRP	OS6900	Advanced
IGMP Multicast Group Configuration Limit	OS6900	Base

Feature	Platform	License
IGMP Relay	OS6900	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS6900	Base
L2 Static Multicast Address	OS6900	Base
PIM / PIM-SSM (Source-Specific Multicast)	OS6900	Advanced
Monitoring/Troubleshooting Feature Support		
DDM - Digital Diagnostic Monitoring	OS6900	Base
Health Statistics	OS6900	Base
Ping and Traceroute	OS6900	Base
Policy Based Mirroring	OS6900	Base
Port Mirroring	OS6900	Base
Port Monitoring	OS6900	Base
Remote Port Mirroring	OS6900	Base
Rmon	OS6900	Base
sFlow	OS6900	Base
Switch Logging and Syslog	OS6900	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS6900	Base
Ethernet Services	OS6900	Base
L2 Control Protocol Tunneling (L2CP)	OS6900	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS6900	Base
Account & Password Policies	OS6900	Base
Admin User Remote Access Restriction Control	OS6900	Base
ARP Defense Optimization	OS6900	Base
ARP Poisoning Detect	OS6900	Base
Authenticated Switch Access	OS6900	Base

Feature	Platform	License
IP DoS Filtering	OS6900	Base
Learned Port Security (LPS)	OS6900	Base
Policy Server Management	OS6900	Base

Existing Hardware/Software Feature Summary - AOS 7.1.1. R01

Feature	Platform	Software Package
Hardware Feature Support		
OmniSwitch 10K Chassis OS10K-CMM OS10K-CFM OS10K-GNI-C48E OS10K-GNI-U48E OS10K-XNI-U32S OS10K-PS-25A OS10K-PS-24D OS10K-Fan-Tray		
Manageability Feature Support		
CLI	OS10K	Base
Ethernet Interfaces	OS10K	Base
ISSU	OS10K	Base
Multiple VRF Routing and Forwarding	OS10K	Base
Network Time Protocol (NTP)	OS10K	Base
Pause Control/Flow Control	OS10K	Base
Remote Access FTP SCP SSH/SFTP Telnet TFTP	OS10K	Base
Smart Continuous Switching Hot Swap Management Module Failover Power Monitoring Redundancy	OS10K	Base
SNMP	OS10K	Base
Software Rollback – Multi-Image/Multi-Config	OS10K	Base
Storm Control	OS10K	Base
Text File Configuration	OS10K	Base

Feature	Platform	Software Package
UDLD	OS10K	Base
USB Support	OS10K	Base
Web-Based Management (WebView)	OS10K	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS10K	Base
802.1Q	OS10K	Base
Configurable Hash Mode	OS10K	Base
Link Aggregation –Static and LACP (802.3ad)	OS10K	Base
Multi-Chassis Link Aggregation	OS10K	Base
Source Learning	OS10K	Base
Spanning Tree • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard	OS10K	Base
VLANs	OS10K	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS10K	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS10K	Base
BGP4 with Graceful Restart	OS10K	Base
DNS Client	OS10K	Base
GRE	OS10K	Base
IP Multicast Routing	OS10K	Base
IP Multicast Switching (IGMP)	OS10K	Base
IP Multicast Switching (Proxying)	OS10K	Base
IP Multinetting	OS10K	Base

Feature	Platform	Software Package
IP Route Map Redistribution	OS10K	Base
IP-IP Tunneling	OS10K	Base
OSPFv2	OS10K	Base
RIPv1/v2	OS10K	Base
Routing Protocol Preference	OS10K	Base
Server Load Balancing	OS10K	Base
VRRPv2	OS10K	Base
IPv6 Feature Support		
BGP4 BGP IPv6 Extensions	OS10K	Base
IPSec IPv6 OSPFv3 RIPng	OS10K	Base
IPv6 Client and/or Server Support	OS10K	Base
IPv6 Multicast Routing	OS10K	Base
IPv6 Multicast Switching (MLD v1/v2)	OS10K	Base
IPv6 Routing	OS10K	Base
IPv6 Scoped Multicast Addresses	OS10K	Base
IPv6 Neighbor Discovery Support	OS10K	Base
OSPFv3	OS10K	Base
RIPng	OS10K	Base
VRRPv3	OS10K	Base
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS10K	Base
Ingress and egress bandwith shaping	OS10K	Base
Policy Based Routing	OS10K	Base
Tri-Color Marking	OS10K	Base
Multicast Feature Support		
DVMRP	OS10K	Base

Feature	Platform	Software Package
IGMP Multicast Group Configuration Limit	OS10K	Base
IGMP Relay	OS10K	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS10K	Base
L2 Static Multicast Address	OS10K	Base
PIM / PIM-SSM (Source-Specific Multicast)	OS10K	Base
Monitoring/Troubleshooting Feature Support		
DDM - Digital Diagnostic Monitoring	OS10K	Base
Health Statistics	OS10K	Base
Ping and Traceroute	OS10K	Base
Policy Based Mirroring	OS10K	Base
Port Mirroring	OS10K	Base
Port Monitoring	OS10K	Base
Remote Port Mirroring	OS10K	Base
Rmon	OS10K	Base
sFlow	OS10K	Base
Switch Logging and Syslog	OS10K	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS10K	Base
Ethernet Services	OS10K	Base
L2 Control Protocol Tunneling (L2CP)	OS10K	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS10K	Base
Account & Password Policies	OS10K	Base
Admin User Remote Access Restriction Control	OS10K	Base
ARP Defense Optimization	OS10K	Base
ARP Poisoning Detect	OS10K	Base

Feature	Platform	Software Package
Authenticated Switch Access	OS10K	Base
IP DoS Filtering	OS10K	Base
Learned Port Security (LPS)	OS10K	Base
Policy Server Management	OS10K	Base

Appendix F: Release Specifications

This appendix is derived from the OmniSwitch AOS user guides. It contains all the specifications tables at the beginning of each chapter in each of user guides of the corresponding release. It is designed to be a single resource to help verify the specifications being documented for AOS releases. The information contained here is duplicated in the Specifications Tables in each user guide.

Swith Management Guide Specifications

Getting Started Specifications

Platforms Supported	OmniSwitch 10K, 6900
Standalone Configuration Files	boot.cfg
Virtual Chassis Configuration Files	vcboot.cfg vcsetup.cfg
Demo License	45-day Demo Advanced license
Image Files	Ros.img (OS10K) Reni.img (OS10K) Tos.img (OS6900)
Validation File	issu_version
ISSU Directory	Any user-defined directory to store the image files
NI Reset Timer	120 minutes
Control LED	Blinks amber during ISSU upgrade

Login Specifications

Platforms Supported	OmniSwitch 10K, 6900
Login Methods	Telnet, SSH, HTTP, SNMP
Number of concurrent Telnet sessions	6
Number of concurrent SSH sessions	8
Number of concurrent HTTP (WebView) sessions	4
Secure Shell public key authentication	Password DSA/RSA Public Key
RFCs Supported for SSHv2	RFC 4253 - SSH Transport Layer Protocol RFC 4418 - UMAC: Message Authentication Code using Universal Hashing

File Management Specifications

Platforms Supported	OmniSwitch 10K, 6900

File Transfer Methods	FTP (v4/v6), SFTP (v4/v6), SCP (v4/v6), TFTP
Client/Server Support	FTP - Client (IPv4 Only) or Server SFTP - Client or Server SCP - Client or Server TFTP – Client
Number of concurrent FTP/ SFTP sessions	4
Configuration Recovery	The **flash/certified** directory holds configurations that are certified as the default start-up files for the switch. They will be used in the event of a non-specified reload.
Default Switch Directory - **/flash**	Contains the **certified, working, switch, network,** and user-defined directories.
File/Directory Name Metrics	255 character maximum. File and directory names are case sensitive.
File/Directory Name Characters	Any valid ASCII character except '/'.
Sub-Directories	Additional user-defined directories created in the **/flash** directory.
Text Editing	Standard Vi standard editor.
System Clock	Set local date, time and time zone, Universal Time Coordinate (UTC), Daylight Savings (DST or summertime).

Managing CMM Directory Content

CMM Specifications

Platforms Supported	**OmniSwitch 10K, 6900**
Size of Flash Memory	2 GB
Maximum Length of File Names	255 Characters
Maximum Length of Directory Names	255 Characters
Maximum Length of System Name	32 Characters
Default Boot Directory	Certified

USB Flash Drive Specifications

Platforms Supported	OmniSwitch 10K, 6900
USB Flash Drive Support	Alcatel-Lucent Certified USB Flash Drive
Automatic Software Upgrade	Supported
Disaster Recovery	Supported OS10K - **Rrescue.img** file required OS6900 - **Trescue.img** file required

Note: The format of the Alcatel-Lucent certfied USB Flash Drive must be FAT32. To avoid file corruption issues the USB Drive should be stopped before removing from a PC. Directory names are case sensitive and must be lower case.

CLI Specifications

Platforms Supported	OmniSwitch 10K, 6900
Configuration Methods	• Online configuration via real-time sessions using CLI commands. • Offline configuration using text file holding CLI commands.
Command Capture Feature	Snapshot feature captures switch configurations in a text file.
User Service Features	• Command Line Editing • Command Prefix Recognition • CLI Prompt Option • Command Help • Keyword Completion • Keyword Abbreviation • Command History • Command Logging • Syntax Error Display • More Command

Configuration File Specifications

Platforms Supported	OmniSwitch 10K, 6900
Creation Methods for Configuration Files	• Create a text file on a word processor and upload it to the switch. • Invoke the switch's snapshot feature to create a text file. • Create a text file using the switch's text editor.
Timer Functions	Files can be applied immediately or by setting a timer on the switch.
Command Capture Feature	Snapshot feature captures switch configurations in a text file.
Error Reporting	Snapshot feature includes error reporting in the text file.
Text Editing on the Switch	Vi standard editor.
Default Error File Limit	1

User Database Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum number of alphanumeric characters in a username	63
Maximum number of alphanumeric characters in a user password	30
Maximum number of local user accounts	50

Switch Security Specifications

Platforms Supported	OmniSwitch 10K, 6900
N/A	N/A

WebView Specifications

Platforms Supported	OmniSwitch 10K, 6900
N/A	N/A

SNMP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported for SNMPv2	1902 through 1907 - SNMPv2c Management Framework 1908 - Coexistence and transitions relating to SNMPv1 and SNMPv2c
RFCs Supported for SNMPv3	2570 – Version 3 of the Internet Standard Network Management Framework 2571 – Architecture for Describing SNMP Management Frameworks 2572 – Message Processing and Dispatching for SNMP 2573 – SNMPv3 Applications 2574 – User-based Security Model (USM) for version 3 SNMP 2575 – View-based Access Control Model (VACM) for SNMP 2576 – Coexistence between SNMP versions
Platforms Supported	OmniSwitch 10K, 6900
SNMPv1, SNMPv2, SNMPv3	The SNMPv3 protocol is ascending compatible with SNMPv1 and v2 and supports all the SNMPv1 and SNMPv2 PDUs
SNMPv1 and SNMPv2 Authentication	Community Strings
SNMPv1, SNMPv2 Encryption	None

SNMPv1 and SNMPv2 Security requests accepted by the switch	Sets and Gets
SNMPv3 Authentication	SHA, MD5
SNMPv3 Encryption	DES
SNMPv3 Security requests accepted by the switch.	Non-authenticated Sets, Non-authenticated Gets and Get-Nexts, Authenticated Sets, Authenticated Gets and Get-Nexts, Encrypted Sets, Encrypted Gets and Get-Nexts

Web Services, CLI Scripting, OpenFlow Specifications

Platforms Supported	OmniSwitch 10K, 6900
Configuration Methods	HTTP/HTTPS Python API
Response Formats	Extensible Markup language (XML) JavaScript Object Notation (JSON)
Maximum Web Services Session	4
Alcatel-Lucent Python Library	consumer.py (Python version 2.X/3.X compatible) Note: This file is available on the Service & Support Website. It is being provided as an example application to help with Web Services familiarization but is not an officially supported part of the Web Services solution.
Internal Python in AOS/Event based CLI Scripting	Python 3
Default Script Run Time Limit 60 Seconds.	60 Seconds

OpenFlow Specifications

Platforms Supported	OmniSwitch 10K, 6900 **Note**: Not supported on OS10K-XNI-U32S module.
Modes Supported	Normal Hybrid (API)
Version Supported	1.0 1.3.1
Maximum Number of logical switches	3
Maximum number of controllers per logical	3

switch	
Maximum number of logical switches in Hybrid mode	1
Support for Virtual Chassis	Supported
OpenFlow 1.0/1.3.1 TCP port	6633

Virtual Chassis Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum number of physical switches in a Virtual Chassis Note: OS10Ks and OS6900s cannot be mixed in a Virtual Chassis Note: Different OS6900 models can be mixed in a Virtual Chassis.	OS10K (2) OS6900 (6)
Valid chassis identifier	OS10K - 1 or 2 OS6900 – 1 through 6
Valid chassis group identifier	0–255
Valid chassis priority	0–255
Maximum number of Virtual Fabric Links	OS10K – 1 OS6900 - 5
Valid Virtual Fabric Link identifier	OS10K – 0 OS6900 – 0 through 4
VFL Supported Port Types	10G or 40G
Valid control VLAN	2-4094
Valid Virtual Chassis protocol hello interval	1-10
Maximum number of member ports per Virtual Fabric Link	16
Licenses Required	Advanced or Demo
OS6900 OK LED	Blinking Green = Master Solid Green = Slave

Note: Distributed MAC Learning Mode is not supported on a Virtual Chassis

Automatic Remote Configuration Specifications

Platforms Supported	OmniSwitch 10K, 6900
DHCP Specifications	DHCP Server required DHCP Client on OmniSwitch - VLAN 1

	- Tagged VLAN 127 (all ports) - LLDP Management VLAN - Automatic LACP (tagged VLAN 127, untagged VLAN 1)
File Servers	TFTP FTP/SFTP
Clients supported	TFTP FTP/SFTP
Instruction file	Maximum length of: • Pathname: 255 characters • Filename: 63 characters
Maximum length of username for FTP/SFTP file server.	15 characters
Maximum DHCP lease tries	6
Unsupported Features	ISSU and IPv6 are not supported. Upgrade of uboot, miniboot, or FPGA files is not supported.
OK LED	Flashing amber during Automatic Remote Configuration process

Automatic Fabric Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Advanced (free 45-day license activated when the switch comes up)
Modes Supported	Standalone or Virtual Chassis
Ports Supported	Any port that is not configured for use by another feature (for example, 802.1q tag, UNP, or Ethernet Services).
IP Protocols Supported for Automatic IP Configuration	OSPFv2, OSPFv3, IS-IS IPv4, IS-IS IPv6

NTP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs supported	1305–Network Time Protocol
NTP Key File Location	**/flash/network**
Platforms Supported	OmniSwitch 10K, 6900
Maximum number of NTP servers	12

Network Configuration Guide Specifications

Ethernet Specifications

IEEE Standards Supported	802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) 802.3u (100BaseTX) 802.3ab (1000BaseT) 802.3z (1000Base-X) 802.3ae (10GBase-X) 802.3ba (40GBase-X) 802.3z (Energy Efficient Ethernet)
Platforms Supported	OmniSwitch 10K, 6900
Ports Supported	Ethernet (10 Mbps) Fast Ethernet (100 Mbps) Gigabit Ethernet (1 Gbps) 10 Gigabit Ethernet (10 Gbps) 40 Gigabit Ethernet (40 Gbps)
Auto Negotiation	Supported
Port Mirroring / Monitoring	Supported
802.1Q Hardware Tagging	Supported
Jumbo Frame Configuration	Supported on 1/10/40 Gigabit Ethernet ports
Maximum Frame Size	1553 bytes (10/100 Mbps) 9216 bytes (1/10/40 Gbps)
Enhanced Port Performance (EPP)	Supported on OS6900 with 10-Gigabit transceivers

UDLD Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum number of UDLD ports per system	Up to maximum physical ports per system

Source Learning Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs supported	2674—Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
Maximum number of learned MAC addresses when centralized MAC source learning mode is enabled	OS10K - 32K Module / 32K Chassis OS6900-X20/X40/T20/T40 - 128K OS6900-Q32 – 228K
Maximum number of learned MAC addresses when distributed MAC source learning mode	OS10K - 32K Module (C48E/U48E/U32S) OS10K – 128K Module (U32E/U16E(L)/U4E/U8E)

is enabled.	OS10K – 256K (Chassis) OS6900 - Not Supported

VLAN Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	2674 - Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
IEEE Standards Supported	802.1Q - Virtual Bridged Local Area Networks 802.1D - Media Access Control Bridges
Maximum VLANs per switch	4094
Maximum Tagged VLANs per Port	4093
Maximum Untagged VLANs per Port	One untagged VLAN (default VLAN) per port.
Maximum VLAN Port Associations (VPA) per switch (Recommended)	OS10K - 20000 OS6900 - 10000
Maximum Spanning Tree VLANs per switch	128 (1x1 mode)

High Availability VLANs Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum high availability VLANs per switch	16
Switch ports eligible for high availability VLAN assignment.	Fixed ports on second-generation Network Interface (NI) modules.
Switch port not eligible for high availability VLAN assignment.	Mirroring ports.

Spanning Tree Specifications

Platforms Supported	OmniSwitch 10K, 6900
IEEE Standards supported	802.1d—Media Access Control (MAC) Bridges 802.1s—Multiple Spanning Trees 802.1w—Rapid Spanning Tree Protocol
Spanning Tree operating modes supported	Flat mode—one spanning tree instance per switch Per-VLAN mode—one spanning tree instance per VLAN
Spanning Tree port eligibility	Fixed ports 802.1Q tagged ports Link aggregate of ports
Maximum VLAN Spanning Tree instances per switch.	252 (per-VLAN mode)
Maximum flat mode Multiple Spanning Tree Instances (MSTI) per switch	16 MSTI, in addition to the Common and Internal Spanning Tree instance (also referred to as MSTI 0).

Loopback Detection Specifications

Platforms Supported	OmniSwitch 10K, 6900
Ports Supported	There is no restriction on the type of ports on which the LBD can be enabled. But it is recommended LBD should be enabled on the edge ports.
Transmisson Timer	Range from 5 to 600 seconds.
Auto-recovery Timer	Range from 30 to 86400 seconds.

Static Link Aggregation Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum number of link aggregation groups	OS10K – 128 OS6900 - 256
Maximum number of links per group supported	OS10K – 8 OS6900 - 16

Dynamic Link Aggregation Specifications

Platforms Supported	OmniSwitch 10K, 6900
IEEE Specifications Supported	802.3ad — Aggregation of Multiple Link Segments
Maximum number of link aggregation groups	OS10K – 128 OS6900 - 256
Maximum number of ports per link aggregate	OS10K – 8 OS6900 - 16

ERP Specifications

ITU-T G.8032 03/2010	Ethernet Ring Protection version 2 (Multi Rings and Ladder networks supported) (Hold off timer, Lockout , Signal degrade SD, RPL Replacement, Forced Switch, Manual Switch, Clear for Manual/Forced Switch, Dual end blocking not supported)
ITU-T Y.1731/IEEE 802.1ag	ERP packet compliant with OAM PDU format for CCM
Supported Platforms	OmniSwitch 10K, 6900
Maximum number of rings per node	64
Maximum number of nodes per ring	16 (recommended)
Maximum number of VLANs per port.	4094

Range for ring ID	1 - 2147483647
Range for remote MEPID	1 - 8191
Range for wait-to-restore timer	1 - 12 minutes
Range for guard timer	1 - 200 centi-seconds

MVRP Specifications

IEEE Standards Supported	IEEE 802.1ak-2007 Amendment 7: Multiple Registration Protocol IEEE 802.1Q-2005 Corrigendum 2008
Platforms Supported	OmniSwitch 10K, 6900
Maximum MVRP VLANs	4094

802.1AB Specifications

Platforms Supported	OmniSwitch 10K, 6900
IEEE Specification	IEEE 802.1AB-2005 Station and Media Access Control Connectivity Discovery
Maximum number of network policies that can be associated with a port	8
Maximum number of network policies that can be configured on the switch	32

IP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	791–Internet Protocol 792–Internet Control Message Protocol 826–An Ethernet Address Resolution Protocol 2784–Generic Routing Encapsulation (GRE) 2890–Key and Sequence Number Extensions to GRE (extensions defined are not supported) 1701–Generic Routing Encapsulation (GRE) 1702–Generic Routing Encapsulation over IPV4 Networks 2003-IP Encapsulation within IP
Maximum router interfaces per system	4094 IP
Maximum router interfaces per VLAN	16
Maximum HW routes	OS10K (C48/U48)- 256K

	OS10K (U32E) – 16K OS10K (U32S) – 12K OS6900-X20/X40/T20/T40 - 16K OS6900-Q32 – 12K
Maximum SW routes per switch	OS10K – 256K OS6900 – 128K
Maximum HW ARP entries per module	OS10K (XNI-U32S) – 8K OS10K (All other modules) - 16K OS6900 (X20/X40) - 8K OS6900 (T20/T40) – 16K OS6900-Q32 – 32K (Note: Mixing an XNI-U32S with other modules in the same chassis reduces the maximum ARP entries to 8K for all modules.)
Maximum number of GRE tunnel interfaces per switch	127
Maximum number of IPIP tunnel interfaces per switch	127
Routing protocols supported over the tunnel interfaces	RIP, OSPF, BGP
Maximum next hops per ECMP entry (static or RIP routes)	16

VRF Specifications

Platforms Supported	OS10K, 6900
OmniSwitch License Requirements	Advanced License required on OmniSwitch 6900 only.
Routing Protocols Supported	Static, IPv4, RIPv2, OSPFv2,BGP4, IS-IS
Maximum number of max profile VRF instances per switch (no low profiles)	64
Maximum number of low profile VRF instances per switch (no high profiles)	445 (OmniSwitch 10K) 128 (OmniSwitch 6900)
Maximum VRF instances per VLAN	1
Maximum OSPF VRF routing instances per switch	16
Maximum RIPv2 VRF routing instances per switch	16
Maximum BGP VRF routing instances per switch	32
SNMP version required for management	SNMPv3

IPv6 Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	1981 Path MTU Discovery for IP version 6 2375 IPv6 Multicast Address Assignments 2460 Internet Protocol, Version 6 (IPv6) Specification 2464 Transmission of IPv6 Packets over Ethernet Networks 2465 Management Information Base for IP Version 6: Textual Conventions and General Group 2466 Management Information Base for IP Version 6: ICMPv6 Group 2711 IPv6 Router Alert Option 3056 Connection of IPv6 Domains via IPv4 Clouds 3484 Default Address Selection for Internet Protocol version 6 (IPv6) 3493 Basic Socket Interface Extensions for IPv6 3542 Advanced Sockets Application Program Interface (API) for IPv6 3587 IPv6 Global Unicast Address Format 3595 Textual Conventions for IPv6 Flow Label 3596 DNS Extensions to Support IP Version 6 4007 IPv6 Scoped Address Architecture 4022 Management Information Base for the Transmission Control Protocol (TCP) 4113 Management Information Base for the User Datagram Protocol (UDP) 4193 Unique Local IPv6 Unicast Addresses 4213 Basic Transition Mechanisms for IPv6 Hosts and Routers 4291 IP Version 6 Addressing Architecture 4294 IPv6 Node Requirements 4443 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification 4861 Neighbor Discovery for IP version 6 (IPv6) 4862 IPv6 Stateless Address Autoconfiguration 5095 Deprecation of Type 0 Routing Headers in IPv6 5453 Reserved IPv6 Interface Identifiers 5722 Handling of Overlapping IPv6 Fragments
Maximum IPv6 interfaces	VLANs- 4096 Configured Tunnels - 255 6to4 Tunnels - 1
Maximum IPv6 global unicast or anycast addressess	10K
Maximum IPv6 global unicast addresses per IPv6 interface	50
Maximum IPv6 addresses assigned via VRRP configuration	1K (per j.m)
Maximum IPv6 hardware routes when there are no IPv4 routes present (includes dynamic and static routes)	OS10K / OS6900 – 256 (prefix >= 65) OS10K (U48/C48) – 8K (prefix <= 64)

	OS10K (U32S) – 6K (prefix <= 64) OS10K (U32E) – 8K (prefix <= 64) OS6900-X/T – 8K (prefix <= 64) OS6900-Q32 – 6K (prefix <= 64) (Note: Exceeding these limits, or having IPv4 routes will result in some traffic being routed in software)
Maximum Number of RIPng Peers	10
Maximum Number of RIPng Interfaces	10
Maximum Number of RIPng Routes	5K
Maximum next hops per ECMP entry (static or RIPng routes)	16

IPsec Specifications

Platforms Supported	OmniSwitch 10K, 6900
IP Version Supported	IPv6
RFCs Supported	4301 - Security Architecture for the Internet Protocol 4302 - IP Authentication Header (AH) 4303 - IP Encapsulating Security Payload (ESP) 4305 - Cryptographic Algorithm Implementation Requirements for ESP and AH 4308 - Cryptographic Suites for IPsec
Encryption Algorithms Supported for ESP	NULL, 3DES-CBC, and AES-CBC
Key lengths supported for Encryption Algorithms	3DES-CBC - 192 bits AES-CBC - 128, 192, or 256 bits
Authentication Algorithms Supported for AH	HMAC-SHA1-96, HMAC-MD5-96, and AES- XCBC-MAC-96
Key lengths supported for Authentication Algorithms	HMAC-MD5 - 128 bits HMAC-SHA1 - 160 bits AES-XCBC-MAC - 128 bits
Master Security Key formats	Hexadecimal (16 bytes) or String (16 characters)
Priority value range for IPsec Policy	1 - 1000
Index value range for IPsec Policy Rule	1 - 10
SPI Range	256 - 999999999
Modes Supported	Transport

RIP Specifications

Platforms Supported	OmniSwitch 10K, 6900

RFCs Supported	RFC 1058–RIP v1 RFC 2453–RIP v2 RFC 1722–RIP v2 Protocol Applicability Statement RFC 1724–RIP v2 MIB Extension
Maximum Number of Interfaces	10
Maximum Number of Peers	100
Maximum Number of Routes	10K
Maximum next hops per ECMP entry	16

BFD Specifications

RFCs Supported	5880—Bidirectional Forwarding Detection 5881—Bidirectional Forwarding Detection for IPv4 and IPv6 (Single Hop) 5882—Generic Application of Bidirectional Forwarding Detection
Platforms Supported	OmniSwitch 10K, 6900
Maximum Number of BFD Sessions	OS6900 (per chassis) - 32 OS6900 (Virtual Chassis) - 100 OS10K (per NI) - 64 OS10K (Chassis/ Virtual Chassis) - 512
Protocols Supported	BGP, OSPF, VRRP Remote Address Tracking only, and Static Routes. IPv6 protocols not supported.
Modes Supported	Asynchronous Echo (Demand Mode not supported)

DHCP Relay Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	0951–Bootstrap Protocol 1534–Interoperation between DHCP and BOOTP 1541–Dynamic Host Configuration Protocol 1542–Clarifications and Extensions for the Bootstrap Protocol 2132–DHCP Options and BOOTP Vendor Extensions 3046–DHCP Relay Agent Information Option, 2001
DHCP Relay Implementation	Global DHCP Per-VLAN DHCP
DHCP Relay Service	BOOTP/DHCP (Bootstrap Protocol/Dynamic Host Configuration Protocol)

UDP Port Numbers	67 for Request 68 for Response
IP addresses supported for each Relay Service	Maximum of 256 IP addresses for each Relay Service.
IP addresses supported for the Per-VLAN service	Maximum of 256 VLAN relay services.
Maximum number of UDP relay services allowed per switch	10
Maximum number of VLANs to which forwarded UDP service port traffic is allowed	256

DHCP Server Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	RFC 2131 - Dynamic Host Configuration Protocol RFC 3315 - Dynamic Host Configuration Protocol for IPv6 RFC 950 - Internet Standard Subnetting Procedure RFC 868 - Time Protocol RFC 1035 - Domain Implementation and Specification RFC 1191- Path MTU Discovery
DHCP Server Implementation	BOOTP/DHCP
UDP Port Numbers	67 for Request and Response (IPv4) 547 for Request (IPv6) 546 for Response (IPv6)
IP address lease allocation mechanisms:	**Static BootP:** IP address is allocated using the BootP configuration when the MAC address of the client is defined. **Static DHCP:** The network administrator assigns an IP address to the client. DHCP conveys the address assigned by the DHCP server to the client. **Dynamic DHCP:** The DHCP server assigns an IP address to a client for a limited period of time or until the client explicitly

	releases the address.
OmniSwitch IPv4 Configuration Files	dhcpd.conf dhcpd.pcy dhcpsrv.db
OmniSwitch IPv6 Configuration Files	dhcpdv6.conf dhcpdv6.pcy dhcpv6srv.db
Maximum number of leases	8000
Maximum lease information file size	375 KB

VRRP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	RFC 3768–Virtual Router Redundancy Protocol RFC 2787–Definitions of Managed Objects for the Virtual Router Redundancy Protocol
Compatible with HSRP	No
Maximum number of VRRPv2 and VRRPv3 virtual routers	255
Maximum number of IP addresses per instance	16

Server Load Balancing Specifications

Platforms Supported	OmniSwitch 10K, 6900
Maximum number of clusters	32
Maximum number of physical servers per cluster	32
Layer-3 classification	Destination IP address QoS policy condition
Layer-2 classification	QoS policy condition
Server health checking	Ping, link checks
High availability support	Hardware-based failover, VRRP, Chassis Management Module (CMM) redundancy
Networking protocols supported	Virtual IP (VIP) addresses
Maximum number of probes on a switch	40

IPMS Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	RFC 1112 — Host Extensions for IP Multicasting RFC 2236 — Internet Group Management Protocol, Version 2 RFC 2710 -- Multicast Listener Discovery (MLD) for IPv6 RFC 2933 — Internet Group Management Protocol MIB RFC 3019 -- IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol RFC 3376 -- Internet Group Management Protocol, Version 3 RFC 3810 — Multicast Listener Discovery Version 2 (MLDv2) for IPv6 RFC 4541 — Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches RFC 4604 — Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast
IGMP Versions Supported	IGMPv1, IGMPv2, IGMPv3
Maximum number of IPv4 multicast flows	OS10K - 4K OS10K – 2K (XNI-U32S) OS6900 (X20/X40) - 2K OS6900 (T20/T40) – 2K OS6900-Q32 – 20K (Note: Mixing an XNI-U32S with other modules in the same chassis reduces the maximum entries to 2K)

IPMSv6 Specifications

RFCs Supported	RFC 2710 — Multicast Listener Discovery for IPv6 RFC 3019 — IPv6 MIB for Multicast Listener Discovery Protocol 3306—Unicast-Prefix-based IPv6 Multicast Addresses RFC 3810 — Multicast Listener Discovery Version 2 for IPv6 RFC 4541 - Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches RFC 4604 - Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicas

Platforms Supported	OmniSwitch 10K, 6900
MLD Versions Supported	MLDv1, MLDv2
MLD Query Interval	1 to 65535 in seconds
MLD Router Timeout	1 to 65535 in seconds
MLD Source Timeout	1 to 65535 in seconds
MLD Query Response Interval	1 to 65535 in milliseconds
MLD Last Member Query Interval	1 to 65535 in milliseconds
Maximum number of IPv6 multicast flows	OS10K - 4K OS10K – 2K (XNI-U32S) OS6900 (X20/X40) - 2K OS6900 (T20/T40) – 2K OS6900-Q32 – 20K (Note: Mixing an XNI-U32S with other modules in the same chassis reduces the maximum entries to 2K)

QoS Specifications

Maximum number of policy rules	8192
Maximum number of policy conditions	8192
Maximum number of policy actions	8192
Maximum number of policy rules per slot	1024 – OS10K-XNI-U32E, OS6900 1280 – OS10K-XNI-U32S 2560 OS6900-Q32 5120 OS10K-GNI-C48E, OS10K-GNI-U48E)
Maximum number of bandwidth policy rules	2560 (OmniSwitch 10K) 512 (OmniSwitch 6900)
Maximum number of validity periods	64
Maximum number of policy services	256
Maximum number of groups (network, MAC, service, port)	2048
Maximum number of group entries	1024 per group (512 per service group)
Maximum number of Class of Service (CoS) queues per port.	8
Queue Set Profiles (QSP)	4
Weighted Random Early Detection profiles (WRP)	1 (OmniSwitch 6900) Not supported on the OmniSwitch 10K

Maximum number of QoS policy lists per switch	32 (includes the default list)
Maximum number of QoS policy lists per Universal Network Profile (UNP)	1

Policy Server Specifications

Platforms Supported	OmniSwitch 10K, 6900
LDAP Policy Servers RFCs Supported	RFC 2251–Lightweight Directory Access Protocol (v3) RFC 3060–Policy Core Information Model—Version 1 Specification
Maximum number of policy servers (supported on the switch)	5
Maximum number of policy servers (supported by PolicyView)	1

UNP Specifications

Platforms Supported	OmniSwitch 6900, 10K
Number of UNPs per switch	4K (Includes static and dynamic profiles)
Number of UNPs users per switch	2K
Authentication Type	MAC and 802.1x
Profile type	VLAN, SPB service, or VXLAN service
UNP port type	Bridge (VLAN-based classification) or access (service-based classification)
UNP classification rules	MAC address, MAC-range, IP address, and VLAN tag
Number of QoS policy lists per switch	32 (includes the default list)
Number of QoS policy lists per User Network Profile	1

Application Fingerprinting Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	N/A
Supported Packet Types	IP (IPv4 and IPv6)
Application signature type	REGEX
AOS provided signatures	Chatting Program, Mail, Networking or IETF Proposal Standard, P2P, Remote Access, VOIP

Authentication Server Specifications

Platforms Supported	OmniSwitch 10K, 6900

RADIUS RFCs Supported	RFC 2865–Remote Authentication Dial In User Service (RADIUS) RFC 2866–RADIUS Accounting RFC 2867–RADIUS Accounting Modifications for Tunnel Protocol Support RFC 2868–RADIUS Attributes for Tunnel Protocol Support RFC 2809–Implementation of L2TP Compulsory Tunneling through RADIUS RFC 2869–RADIUS Extensions RFC 2548–Microsoft Vendor-specific RADIUS Attributes RFC 2882–Network Access Servers Requirements: Extended RADIUS Practices
TACACS+ RFCs Supported	RFC 1492–An Access Control Protocol
LDAP RFCs Supported	RFC 1789–Connectionless Lightweight X.5000 Directory Access Protocol RFC 2247–Using Domains in LDAP/X.500 Distinguished Names RFC 2251–Lightweight Directory Access Protocol (v3) RFC 2252–Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions RFC 2253–Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names RFC 2254–The String Representation of LDAP Search Filters RFC 2256–A Summary of the X.500(96) User Schema for Use with LDAPv3
Other RFCs	RFC 2574–User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) RFC 2924–Accounting Attributes and Record Formats RFC 2975–Introduction to Accounting Management RFC 2989–Criteria for Evaluating AAA Protocols for Network Access
Maximum number of authentication servers in single authority mode	8
Maximum number of authentication servers in multiple authority mode	8
Maximum number of servers per Authenticated Switch Access type	8

Port Mapping Specifications

Platforms Supported	OmniSwitch 10K, 6900
Ports Supported	Ethernet (10 Mbps) Fast Ethernet (100 Mbps) Gigabit Ethernet (1 Gbps) 10 Gigabit Ethernet (10 Gbps) 40 Gigabit Ethernet (40 Gbps)
Port Mapping Sessions	8

Learned Port Security Specifications

Platforms Supported	OmniSwitch 10K, 6900
Ports eligible for Learned Port Security	Fixed and 802.1Q tagged
Ports not eligible for Learned Port Security	Link aggregate ports. 802.1Q (trunked) link aggregate ports.
Minimum number of learned MAC addresses allowed per LPS port	1
Maximum number of learned MAC addresses allowed per LPS port	1000
Maximum number of filtered MAC addresses allowed per LPS port	100
Maximum number of configurable MAC address ranges per LPS port	1

Diagnosing Switch Problems Specifications

Port Mirroring Specifications

Platforms Supported	OmniSwitch 10K, 6900
Ports Supported	Ethernet (10 Mbps) Fast Ethernet (100 Mbps) Gigabit Ethernet (1 Gbps) 10 Gigabit Ethernet (10 Gbps) 40 Gigabit Ethernet (40 Gbps)
Mirroring Sessions Supported	OmniSwitch 10K - 2 (OS10-XNI-U32 supports 1 session) OmniSwitch 6900 - 2
Combined Mirroring/Monitoring Sessions per Chassis	OmniSwitch 10K – 3 OmniSwitch 6900 - 2
N-to-1 Mirroring Supported	128 to 1
Number of RPMIR VLANs per session	1

Port Monitoring Specifications

Platforms Supported	OmniSwitch 10K, 6900
Ports Supported	Ethernet (10 Mbps) Fast Ethernet (100 Mbps) Gigabit Ethernet (1 Gbps) 10 Gigabit Ethernet (10 Gbps) 40 Gigabit Ethernet (40 Gbps)
Monitoring Sessions Supported	OmniSwitch 10K – 1 OmniSwitch 6900 - 1
Combined Mirroring/Monitoring Sessions per Chassis	OmniSwitch 10K - 3 OmniSwitch 6900 - 2
File Type Supported	ENC file format (Network General Sniffer Network Analyzer Format)

sFlow Specifications

RFCs Supported	3176 - sFlow Management Information Base
Platforms Supported	OmniSwitch 10K, 6900
Receiver/Sampler/Polling Instances	2
Sampling	length of packet type of frame source and destination MACs source and destination VLANs source and destination priorities source and destination IP addressessource and destination ports tcp flags and tos
Polling	In octets Out octets Number of Rx Unicast packets Number of Tx Unicast packets Number of Rx Multicast packets Number of Tx Multicast packets Number of Rx Broadcast packets Number of Tx Broadcast packets In Errors Out Errors

RMON Specifications

RFCs Supported	2819 - Remote Network Monitoring Management

	Information Base
Platforms Supported	OmniSwitch 10K, 6900
RMON Functionality Supported	Basic RMON 4 group implementation –Ethernet Statistics group –History (Control and Statistics) group –Alarms group –Events group
RMON Functionality Not Supported	RMON 10 group* RMON2* –Host group –HostTopN group –Matrix group –Filter group –Packet Capture group (*An external RMON probe that includes RMON 10 group and RMON2 be used where full RMON probe functionality is required.)
Flavor (Probe Type)	Ethernet/History/Alarm
Status	Active/Creating/Inactive
History Control Interval (seconds)	1 to 3600
History Sample Index Range	1 to 65535
Alarm Interval (seconds)	1 to 2147483647
Alarm Startup Alarm	Rising Alarm/Falling Alarm/ RisingOrFalling Alarm
Alarm Sample Type	Delta Value/Absolute
RMON Traps Supported	RisingAlarm/FallingAlarm These traps are generated whenever an Alarm entry crosses either its Rising Threshold or its Falling Threshold and generates an event configured for sending SNMP traps.

Switch Health Specifications

Platforms Supported	OmniSwitch 10K, 6900
Health Functionality Supported	– Switch level CPU Utilization Statistics (percentage); – Switch/module/port level Input Utilization Statistics (percentage); – Switch/module/port level Input/Output Utilization Statistics (percentage); – Switch level Memory Utilization Statistics (percentage);

	– Device level (e.g., Chassis/CMM) Temperature Statistics (Celsius).
Monitored Resource Utilization Levels	–Most recent utilization level; –Average utilization level during last minute; –Average utilization level during last hour; –Maximum utilization level during last hour.
Resource Utilization Raw Sample Values	Saved for previous 60 seconds.
Resource Utilization Current Sample Values	Stored.
Resource Utilization Maximum Utilization Value	Calculated for previous 60 seconds and stored.
Utilization Value = 0	Indicates that none of the resources were measured for the period.
Utilization Value = 1	Indicates that a non-zero amount of the resource (less than 2%) was measured for the period.
Percentage Utilization Values	Calculated based on Resource Measured During Period/Total Capacity.
Resource Threshold Levels	Apply automatically across all levels of switch (switch/module/port).
Rising Threshold Crossing	A Resource Threshold was exceeded by its corresponding utilization value in the current cycle.
Falling Threshold Crossing	A Resource Threshold was exceeded by its corresponding utilization value in the previous cycle, but is not exceeded in the current cycle.
Threshold Crossing Traps Supported	Device, module, port-level threshold crossings.

VLAN Stacking Specifications

Platforms Supported	OmniSwitch 10K, 6900
IEEE Standards Supported	IEEE 802.1Q, 2003 Edition, IEEE Standards for Local and metropolitan area networks—Virtual Bridged Local Area Networks P802.1ad/D6.0 (C/LM) Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment 4: Provider Bridges
Maximum number of Services	4K
Maximum number of SVLANs	4K
Maximum number of SAPs	8K
Maximum number of SAP Profiles	8K (1K if profiles assign priority or bandwidth)
Maximum number of SAP profile VLAN translation or double tagging rules	8K (4K on OS10K XNI-U32 module)

Maximum number of customer VLANs (CVLANs) associated with a SAP	4K
Maximum number of service-to-SAP associations	1K

Switch Logging Specifications

Platforms Supported	OmniSwitch 10K, 6900
Functionality Supported	High-level event logging mechanism that forwards requests from applications to enabled logging devices.
Functionality Not Supported	Not intended for debugging individual hardware applications.
Number of Syslog Servers Supported	12
Logging Devices	Flash Memory/Console/IP Address
Application ID Levels Supported	IDLE (255), DIAG (0), IPC-DIAG (1), QDRIVER (2), QDISPATCHER (3), IPC-LINK (4), NI- SUPERVISION (5), INTERFACE (6), 802.1Q (7), VLAN (8), GM (9), BRIDGE (10), STP (11), LINKAGG (12), QOS (13), RSVP (14), IP (15), IPMS (17), AMAP (18), GMAP (19), SLB(25), AAA (20), IPC-MON (21), IP-HELPER (22), PMM (23), MODULE (24), EIPC (26), CHASSIS (64), PORT-MGR (65), CONFIG (66), CLI (67), SNMP (68), WEB (69), MIPGW (70), SESSION (71), TRAP (72), POLICY (73), DRC (74), SYSTEM (75), HEALTH (76), NAN-DRIVER (78), RMON (79), TELENET (80), PSM (81), FTP (82), SNMI (83), DISTRIB (84), EPIL0GUE (85), LDAP (86), NOSNMP (87), SSL (88), DBGGW (89), LANPOWER (108)
Severity Levels/Types Supported	2 (Alarm - highest severity), 3 (Error), 4 (Alert), 5 (Warning) 6 (Info - default), 7 (Debug 1), 8 (Debug 2), 9 (Debug 3 - lowest severity)

Ethernet OAM Specifications

Platforms Supported	OmniSwitch 10K, 6900
Standards Supported	IEEE 802.1ag Version 8.1–Connectivity Fault Management IEEE 802.1D–Media Access Control (MAC) Bridges IEEE 802.1Q–Virtual Bridged Local Area Networks ITU-T Y.1731–OAM Functions and Mechanisms for Ethernet-Based Networks

Maximum Maintenance Domains (MD) per Bridge	8
Maximum Maintenance Associations (MA) per Bridge	128
Maximum Maintenance End Points (MEP) per Bridge	256
Maximum MEP CMM Database Size	1K

Service Assurance Specifications

Platforms Supported	OmniSwitch 10K, 6900
Standards Supported	N/A

Advanced Routing Guide Specifications

OSPF Specifications

Platforms supported	OmniSwitch 10K, 6900
RFCs supported	1370—Applicability Statement for OSPF 1850—OSPF Version 2 Management Information Base 2328—OSPF Version 2 2370—The OSPF Opaque LSA Option 3101—The OSPF Not-So-Stubby Area (NSSA) Option 3623—Graceful OSPF Restart
Maximum number of areas	10
Maximum number of interfaces per router	128
Maximum number of interfaces per area	100
Maximum number of Link State Database entries	100K
Maximum number of neighbors per router	254
Maximum number of neighbors per area	254
Maximum number of routes	OS10K - 64K OS6900 – 32K (Depending on the number of interfaces/ neighbors, this value may vary.)
License Requirements	Advanced License required on OmniSwitch 6900 only.

OSPFv3 Specifications

Platforms supported	OmniSwitch 10K, 6900
RFCs supported	RFC 1826—IP Authentication Header RFC 1827—IP Encapsulating Security Payload RFC 2553—Basic Socket Interface Extensions for IPv6 RFC 2373—IPv6 Addressing Architecture RFC 2374—An IPv6 Aggregatable Global Unicast Address Format RFC 2460—IPv6 base specification RFC 2470—OSPF for IPv6
Maximum number of areas	5
Maximum number of interfaces per router	20
Maximum number of interfaces per area	16
Maximum number of Link State Database entries per router	20K

Maximum number of neighbors per router	128
Maximum number of neighbors per area	16
Maximum number of routes per router	10K (Depending on the number of interfaces/neighbors, this value may vary.)
License Requirements	Advanced License required on OmniSwitch 6900 only.

ISIS Specifications

Platforms supported	OmniSwitch 10K, 6900
RFCs supported	1142-OSI IS-IS Intra-domain Routing Protocol 1195-OSI IS-IS for Routing in TCP/IP and Dual Environments 3373-Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Pointto-Point Adjacencies 3567-Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication 2966-Prefix Distribution with two-level IS-IS (Route Leaking) support 2763-Dynamic Host name exchange support 3719-Recommendations for Interoperable Networks using IS-IS 3787-Recommendations for Interoperable IP Networks using IS-IS draft-ietf-isis-igp-p2p-over-lan-05.txt-Point-topoint operation over LAN in link-state rout ing protocols 5308 - IS-IS support for IPv6 (Routing IPv6 with IS-IS)
Maximum number of areas (per router)	3
Maximum number of L1 adjacencies per interface (per router)	70
Maximum number of L2 adjacencies per interface (per router)	70
Maximum number of IS-IS interfaces (per router)	70
Maximum number of Link State Packet entries (per adjacency)	255
Maximum number of IS-IS routes	24000
Maximum number of IS-IS L1 routes	12000
Maximum number of IS-IS L2 routes	12000

License Requirements	Advanced License required on OmniSwitch 6900 only.

BGP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	1771/4271–A Border Gateway Protocol 4 (BGP-4) 2439–BGP Route Flap Damping 3392/5492–Capabilities Advertisement with BGP-4 2385–Protection of BGP Sessions via the TCP MD5 Signature Option 1997–BGP Communities Attribute 4456–BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) 3065–Autonomous System Confederations for BGP 4273–Definitions of Managed Objects for BGP-4 4486–Subcodes for BGP Cease Notification 4760–Multiprotocol Extensions for BGP-4 2545–Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing 2918 - Route Refresh Capability for BGP-4 4724 - Graceful Restart Mechanism for BGP 6793 – BGP 4-Octet ASN 5668- 4-Octet AS Specific BGP Extended Community
BGP Attributes Supported	Origin, AS Path, Next Hop (IPv4), MED, Local Preference, Atomic Aggregate, Aggregator (IPv4), Community, Originator ID, Cluster List, Multiprotocol Reachable NLRI (IPv6), Multiprotocol Unreachable NLRI (IPv6), AS4 Path, AS4 Aggregator (IPv4), AS, Specific Extended Community.
Maximum number of peers	512
Maximum number of networks	4K
Maximum number of aggregation addresses	2K
Maximum number of routes	OS10K - 256K OS6900 – 128K
Maximum number of policies	1K
License Requirements	Advanced License required on OmniSwitch 6900 only.

Multicast Boundary Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	2365—Administratively Scoped IP Multicast 5132 - IP Multicast MIB

Valid Scoped Address Range	239.0.0.0 to 239.255.255.255
License Requirements	Advanced License required on OmniSwitch 6900 only.

DVMRP Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs supported	1075 - Distance Vector Multicast Routing Protocol, Version1 4087—IP Tunnel MIB draft-ietf-idmr-dvmrp-v3-09.txt - Distance Vector Multicast Routing Protocol, Version 3 2715—Interoperability Rules for Multicast Routing Protocols
DVMRP version supported	DVMRPv3.255
DVMRP attributes supported	Reverse Path Multicasting, Neighbor Discovery, Multicast Source Location, Route Report Messages, Distance metrics, Dependent Downstream Routers, Poison Reverse, Pruning, Grafting, DVMRP Tunnels
DVMRP timers supported	Flash update interval, Graft retransmissions, Neighbor probe interval, Neighbor timeout, Prune lifetime, Prune retransmission, Route report interval, Route hold-down, Route expiration timeout
Maximum number of interfaces	384 **Note**: Maximum 384 combined Multicast Interfaces between PIMv4, PIMv6 and DVMRP
Multicast protocols per interface	1 (PIM and DVMRP cannot be enabled on the same interface)
License Requirements	Advanced License required on OmniSwitch 6900 only.

PIM Specifications

Platforms supported	OmniSwitch 10K, 6900
RFCs supported	2365 - Administratively Scoped IP Multicast 4601—Protocol Independent Multicast-Sparse Mode (PIM-SM) Protocol Specification 4007 - IPv6 Scoped IP Multicast 5060 - Protocol Independent Multicast MIB 5132 —IP Multicast MIB 3569—An Overview of Source-Specific Multicast (SSM) 3973—Protocol Independent Multicast-Dense Mode (PIM-DM) 5059 - Bootstrap Router (BSR) Mechanism for PIM 5240 - Protocol Independent Multicast (PIM) Bootstrap

	Router MIB 2715—Interoperability Rules for Multicast Routing Protocols
PIM-SM version supported	PIM-SMv2
PIM attributes supported	Shared trees (also referred to as RP trees), Designated Routers (DRs), Bootstrap Routers (BSRs), Candidate Bootstrap Routers (C-BSRs), Rendezvous Points (RPs) (applicable only for PIM-SM), Candidate Rendezvous Points (C-RPs)
PIM timers supported	C-RP expiry, C-RP holdtime, C-RP advertisement, Join/Prune, Probe, Register suppression, Hello, Expiry, Assert, Neighbor liveness
Maximum PIM interfaces	384 **Note**: Maximum 384 combined Multicast Interfaces between PIMv4, PIMv6 and DVMRP
Maximum Rendezvous Point (RP)	100
Maximum Bootstrap Routers (BSRs)	1
Multicast Protocols per Interface	1 (PIM and DVMRP cannot be enabled on the same IP interface)
Valid SSM IPv4 Address Ranges	232.0.0.0 to 232.255.255.255
Valid SSM IPv6 Address Ranges	FF3x::/32
License Requirements	Advanced License required on OmniSwitch 6900 only

Multicast Border Router Specifications

Platforms Supported	OmniSwitch 10K, 6900
RFCs Supported	4601—Protocol Independent Multicast-Sparse Mode (PIM-SM) Protocol Specification 3973—Protocol Independent Multicast-Dense Mode (PIM-DM) 2715—Interoperability Rules for Multicast Routing Protocols draft-ietf-idmr-dvmrp-v3-09.txt - Distance Vector Multicast Routing Protocol, Version 3
MBR Interoperability	DVMRP interoperability with IPv4 PIM (PIM-SM and PIM-DM only).
OmniSwitch License Requirements	Advanced License required on OmniSwitch 6900 only.

Data Center Switching Guide Specifications

DCB Specifications

Platforms Supported	OmniSwitch 6900 and the following OmniSwitch 10K modules: • OS10K-QNI-U8 (8 x 40G) • OS10K-QNI-U4 (4 x 40G) • OS10K-XNI-U32E (32 x 10G) • OS10K-XNI-U16E (16 x 10G) • OS10K-XNI-U16L (8 x 10G, 8 x 1G)
OmniSwitch Software License	Data Center
IEEE Standards supported	802.1Qbb—Priority-based Flow Control 802.1Qaz D2.5—Enhanced Transmission Selection 802.1Qaz D2.5—Data Center Bridging Exchange 802.1Q-REV/D1.5—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks
Maximum number of DCB profiles	128 profiles: • Profiles 1–11 are predefined, with profile 8 serving as the default profile for all ports. • Profiles 12–128 are reserved for user-defined (custom) profiles.
Maximum number of lossless queues (priorities)	110 total per switch (OmniSwitch 6900) 8 per-port (OmniSwitch 10K)
DCB TLVs supported	ETS Configuration ETS Recommendation PFC Configuration (Application Priority TLV not supported)

Shortest Path Bridging Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Advanced
IEEE Standards supported	802.1aq/D3.6: Draft February 10, 2011—Virtual Bridged Local Area Networks-Amendment 9: Shortest Path Bridging 802.1ah/D4.2: DRAFT March 26, 2008— Virtual Bridged

	Local Area Networks–Amendment 6: Provider Backbone Bridging
IETF Internet-Drafts Supported	draft-ietf-isis-ieee-aq-05.txt—ISIS Extensions Supporting IEEE 802.1aq Shortest Path Bridging IETF draft—IP/IPVPN services with IEEE 802.1aq SPBB networks IETF draft—IP/IPVPN services with IEEE 802.1aq SPB networks
SPB Mode Supported	SPB-M (MAC-in-MAC)
IP over SPB-M	IPv4 (VPN-Lite and L3 VPN) VRF-to-ISID mapping (one-to-one, many-to-one)
Maximum number of ISIS-SPB instances per switch.	1
Maximum number of BVLANs per switch	4
Number of equal cost tree (ECT) algorithms supported.	16
Maximum number of service instance identifiers (I-SIDs) per switch	OS6900-Q32 – 8K All other models - 1K
Maximum number of VLANs or SVLANs per I-SID	4K
Maximum number of SAPS	OS10K - 8K OS6900 –X40/X40 – 4K OS6900-T20/T40 – 4K (VC Mode) OS6900-T20/T40 - 8K (Standalone Mode) OS600-Q32 (8K) Note: In a mixed VC the maximum is 4K.
Maximum Transmission Unit (MTU) size for SPB services.	9K (not configurable at this time)

FIP Snooping Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Data Center
INCITS Standards Supported	T11—Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00 June 4, 2009 FC-BB-5 Annex C: Increasing FC-BB_E Robustness Using Access Control Lists T11—Switch Fabric - 5 (FC-SW-5) Rev 8.5 June 3,

	2009
Maximum number of FIP Snooping Sessions	128
Port types supported	10G or faster Ethernet with DCB profile and DCBx enabled with PFC/ETS active (ports and link aggregates)

FCoE Gateway Specifications

Platforms Supported	OmniSwitch 6900 (7.3.3)
OmniSwitch Software License	Data Center
INCITS Standards Supported	• FC-PI-4 Fibre Channel T11/08-138v1 • FC-PI-5 Fibre Channel T11 2118-D/Rev 6.10 • FC-BB-5 Backbone 5 T11/1871-D • FC-BB-6 Backbone 6 T11/2159-D (CNA switching only)
Fibre Channel functionality supported	• FCoE transit bridge - FCoE tunneling of encapsulated FC frames - FCoE initialization protocol (FIP) snooping • FCoE/FC gateway switch - N_Port proxy (NPIV) - F_Port proxy (Reverse-NPIV) - E_Port proxy (E2E-tunnel)
Supported Port types	• Fibre Channel for NPIV gateway—OS-XNI-U12E module with SFP-FC-SR transceiver • Ethernet for FCoE/FIP snooping—10G or faster with DCB profile, DCBx enabled with PFC/ETS active (ports and link aggregates)
OmniSwitch 64-bit World Wide Node Name (WWNN)	10:00:xx:xx:xx:xx:xx:xx (xx = switch MAC address)
OmniSwitch 64-bit World Wide Port Name (WWPN) for each Fibre Channel port	10:00:xx:xx:xx:xx:xx:xx (xx = port MAC address)
VSAN–FC port associations Multiple FC port assignments per VSAN allowed.	Only one VSAN assignment per FC port allowed.
VSAN–FCoE VLAN mapping	One-to-One
VSAN scalability per switch	Based on the number of FC ports (for example, if switch has 12 FC ports, then 12 VSANs; one for each FC port). Note that an FC port configured as an E2E tunnel endpoint does not use up a VSAN assignment.
Maximum number of VSANs per network	4094
E2E tunnel scalability	One tunnel termination per FC port up to the number

	of available FC ports on the switch or virtual chassis.
MTU size supported for SANs	2180
Load Balancing	NP_Port load balancing only: • Dynamic • Dynamic-reorder • ENode-based • Static

Virtual Machine Classification Specifications

UNP (vNP) Specifications

Platforms Supported	OmniSwitch 10K, 6900
Number of UNPs per switch	4K (includes static and dynamic profiles).
Number of UNP users per switch	2K
Authentication type	MAC-based authentication
Profile type	VLAN or Shortest Path Bridging (SPB)
UNP port type	Bridge (VLAN-based classification) or access (service-based classification)
UNP classification rules	MAC address, MAC-range, IP address, and VLAN tag
Number of QoS policy lists per switch	32 (includes the default list)
Number of QoS policy lists per UNP	1

EVB Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Data Center
IEEE Standards Supported	P802.1Qbg Standard Draft, Revision D2.2. February 18, 2012—Virtual Bridged Local Area Networks–Amendment 21: Edge Virtual Bridging
EVB mode	Bridging (virtual machines request the required CVLAN ID tag)
Edge Relay (ER) support	Single ER per switch port. The ER can operate as a Virtual Ethernet Port Aggregator (VEPA) or as a Virtual Ethernet Bridge (VEB).

VXLAN Specifications

Platforms Supported	OmniSwitch 6900-Q32
OmniSwitch Software License	Advanced
RFCs Supported	7348 —VXLAN: A Framework for Overlaying Layer 2 Virtualized Networks over Layer 3 Networks.
VXLAN segments (L2 overlay networks)	16 million
VXLAN service instances	8K
VXLAN Tunnel End Points in a VXLAN network.	500
VXLAN UDP destination ports	1 (default UDP port number is 4789).
VXLAN Service Access Points (SAP)	8K (per device or per Virtual Chassis)
VXLAN SAPs with a VLAN ID range	8 SAPs per service access port
Service access ports with SAPs that contain a VLAN ID range	255
VXLAN Network IDs (VNIs)	4K
Multicast Groups	500
Multicast protocol supported	Bidirectional PIM (BIDIR-PIM)

VXLAN Snooping Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	No software license required
RFCs Supported	7348—VXLAN: A Framework for Overlaying Layer 2 Virtualized Networks over Layer 3 Networks.
Packet Sampling Rate	1K packets-per-second on each module