

Release Notes

OmniSwitch 10K, 6900

Release 7.3.1.R01

These release notes accompany release 7.3.1.R01 software which is supported on the OmniSwitch OS10K and OmniSwitch 6900 platforms. These release notes provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

ATTENTION: Please refer to the **7.3.1.R01 Prerequisite** section for important release specific information prior to upgrading including specific build information for hardware support.

Contents

Related Documentation	3
System Requirements	4
AOS Release 7.3.1.R01 Prerequisites	6
New Hardware Support	6
New Software Features and Enhancements	8
Limited Availability Features and Enhancements.....	17
Existing Hardware Support.....	18
Existing Software Support	22
SNMP Traps.....	59
Unsupported Software Features.....	69
Unsupported CLI Commands	69
Open Problem Reports and Feature Exceptions.....	70
Hot Swap/Redundancy Feature Guidelines.....	74
Technical Support.....	75
Upgrading an OmniSwitch to 7.3.1.R01	76
7.3.1.R01 MC-LAG Upgrade Overview.....	85

Related Documentation

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

OmniSwitch 10K Series Getting Started Guide

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

OmniSwitch 10K Series Hardware User Guide

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

OmniSwitch 6900 Series Getting Started Guide

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

OmniSwitch 6900 Series Hardware User Guide

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

OmniSwitch AOS Release 7 CLI Reference Guide

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

OmniSwitch AOS Release 7 Network Configuration Guide

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

OmniSwitch AOS Release 7 Switch Management Guide

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

OmniSwitch AOS Release 7 Advanced Routing Configuration Guide

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

OmniSwitch AOS Release 7 Data Center Switching Guide

Includes an 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 7.3.1.R01 requires 2GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.

OmniSwitch 10K Series Release 7.3.1.R01 requires 4GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.

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

UBoot and FPGA Requirements

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

Switches not running the minimum version required should upgrade to the latest UBoot or FPGA that is available with the 7.3.1.R01 AOS software available from Service & Support.

- Newly supported OS10K modules released in 7.3.1 listed in the **New Hardware Support** section will be factory shipped with the correct Uboot/FPGA. They do not need to be upgraded and should not be downgraded.
- If upgrading from 7.2.1.R02 the Uboot and FPGA should already be at the correct versions listed below. If upgrading from a release prior to 7.2.1.R02 upgrading the Uboot and FPGA according to the table below is required.
- A separate file containing the Uboot and FPGA upgrade files is available from Service & Support.
- Please refer to the **Upgrade Instructions** section at the end of these Release Notes for step-by-step instructions on upgrading your switch to support 7.3.1.R01.

OmniSwitch 10K

Release	UBoot CMM/NI	FPGA/CPLD CMM	FPGA/CPLD NI
7.3.1.632.R01 (GA)	7.2.1.266.R02	No Upgrade Required	No Upgrade Required

OmniSwitch 6900

Release	UBoot CMM	FPGA/CPLD CMM	FPGA/CPLD Expansion Module
7.3.1.632.R01 (GA)	7.2.1.266.R02	1.3.0 1.2.0	No Upgrade Required

The examples below show how to view the current UBoot and FPGA version on an OmniSwitch.

```
OS10K-> show hardware-info
```

```
CPU Manufacture      : Freescale Semiconductor
CPU Model            : MPC 8572
Compact Flash Manufacturer : CF 2GB
Compact Flash size   : 2097930240 bytes
```

```
RAM Manufacturer      : Other
RAM size             : 3998816 kB
CPM FPGA version     : 2.0
U-Boot Version      : 7.2.1.117.R01 <- UBoot update required
CFMs Present         : 1,2,3,4
Power Supplies Present : 1,2,3,4,-,-,-
Fan Trays Present    : 1,2
NIs Present          : 1,2,3,4,5,6,-,8
```

OS6900-> show hardware-info

```
CPU Manufacture      : Freescale Semiconductor
CPU Model            : MPC 8572
Compact Flash Manufacturer : CF 2GB
Compact Flash size   : 2097930240 bytes
RAM Manufacturer     : Other
RAM size             : 2071912 kB
FPGA 1 version     : 1.0.0 <- FPGA update required
FPGA 2 version     : 1.0.0 <- FPGA update required
U-Boot Version    : 7.2.1.117.R01 <-UBoot update required
Power Supplies Present : 1
NIs Present          : 1,2
```

AOS Release 7.3.1.R01 Prerequisites

Prior to upgrading to AOS Release 7.3.1.R01 please note the following:

Upgrade Instructions – 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 7.3.1.R01.

MC-LAG Upgrade Instructions – If running in an MC-LAG environment the upgrade procedure can be optimized to reduce downtime. Refer to the [MC-LAG Upgrade Instructions](#) for step-by-step instructions.

Data Center License with existing QSet profiles – QSet profiles and DCB profiles are mutually exclusive. If the OmniSwitch Data Center software license is installed, then DCB profiles are used. If this license is not installed, then QSet profiles are used. When a Data Center license is applied any existing QSet profiles will be replaced with DCB profile 8.

In-Service Software Upgrade (ISSU) - Upgrading using ISSU from 7.3.1.519.R01 to 7.3.1.632.R01 is not supported.

Note: The SAP/SPB-M Services, Customer Domains, Dynamic SAP, and Virtual Chassis features are being 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

New Hardware Support

OS10K-XNI-U16L

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

OS10K-XNI-U16E

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

OS10K-XNI-U32E

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

OS10K-QNI-U4E

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

OS10K-QNI-U8E

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

QSFP-40G-LR Transceiver

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

SFP-10G-24DWD80 Transceiver

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

SFP-10G-GIG-SR Transceiver

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

New Software Features and Enhancements

The following software features are being introduced with the 7.3.1.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.

New Software Feature Summary Table

Feature	Platform	License
Data Center Feature Support		
Shortest Path Bridging (SPB)	OS10K/6900	Advanced
Data Center Bridging <ul style="list-style-type: none"> • DCBX • ETS • PFC 	OS10K/6900 OS10K/6900 OS10K/6900	Data Center Data Center Data Center
Edge Virtual Bridging (EVB)	OS10K/6900	Data Center
Virtual Network Profiles <ul style="list-style-type: none"> • SAP/SPB-M Services • Customer Domains (Multi-tenancy) • Dynamic SAP • UNP over MC-LAG on OS10K 	OS10K/6900 OS10K/6900 OS10K/6900 OS10K/6900	Base Base Base Base
Layer 2 Feature Support		
Ethernet Ring Protection v2 (ERPv2)	OS10K/6900	Base
Layer 3 Feature Support		
VRF Management	OS10K/6900	Base
VRF Route Leak	OS10K/6900	Base
Management Feature Support		
Virtual Chassis	OS10K/6900	Advanced
SFP+ Line Diags & Enhanced Port Performance (EPP)	OS10K/6900	Base
License Management	OS10K/6900	Base
Ethernet OAM <ul style="list-style-type: none"> • 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 are being 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.

New Software Features and Enhancements Descriptions

Data Center Feature Descriptions

IEEE 802.1q Shortest Path Bridging

IEEE 802.1aq Shortest Path Bridging (SPB) is designed to expand Layer 2 Ethernet domains and provide multi-path and resiliency capabilities by implementing frame forwarding on the shortest path between any two bridges in an Ethernet network. 802.1aq incorporates two Ethernet encapsulating data points, 802.1ad Provider Bridges (PB) Q-in-Q (Currently not supported) and 802.1ah Provider Backbone Bridges (PBB) MAC-in-MAC (SPB-M)

The Alcatel-Lucent OmniSwitch supports SPB MAC (SPB-M), as defined in the IEEE 802.1aq standard. SPB-M uses the Provider Backbone Bridge (PBB) network model to encapsulate (using IEEE 802.1ah headers) and tunnel customer traffic through the network backbone. The shortest path trees upon which the PBB network infrastructure operates are determined using a version of the Intermediate System-to-Intermediate System (IS-IS) link state protocol that supports TLV extensions for SPB (ISIS-SPB).

Below are some of the key features of the OmniSwitch SPB implementation:

- Multiple shortest paths (Up to 16 paths)
- Deterministic and predictable forwarding
- Compatible with all 802.1, Data Center Bridging protocols, and OA&M
- Allow implementation of free-form POD/MESH topologies
- Fast sub-second convergence

Note: The OS10K-XNI-U32S supports the Backbone Core Bridge (BCB) functionality only.

Data Center Bridging

Data Center Bridging (DCB) is a group of protocols that enhance the current 802.1q Ethernet standard to address the need for a loss-less Ethernet infrastructure. DCB provides the loss-less infrastructure which will enable FCoE and augment iSCSI storage traffic but isn't limited to those applications. The DCB protocols below are supported on the OS10K-XNI-U16L, OS10K-XNI-U16E, OS10K-XNI-U32E, OS10K-QNI-U4, and OS10K-QNI-U8. (bulleted list)

- Priority Flow Control (PFC) 802.1qbb - The primary method to make Ethernet as loss-less as possible is to enforce end to end flow control by using enhanced PAUSE mechanisms to signal backpressure. The current Ethernet PAUSE mechanisms are not sufficient for a true converged Ethernet loss-less network, carrying IP and loss-less traffic, because it would be undesirable for the PAUSE function to impact both types of traffic. To achieve backpressure, a method of segregating the two traffic types into separate priorities and queues (or even different vlans) has been adopted, along with a new PAUSE and internal flow control method. This is called Priority Flow Control (PFC). PFC is done at the link level through priority based PAUSE messages sent 'upstream' between two ports that are PFC enabled.
- Data Center Branch eXchange (DCBX) 802.1qaz- DCBX is an extension to LLDP facilitate communication between switches and hosts in the data center network their capabilities and requirements for datacenter Ethernet functions. To build a fully converged and loss-less capable data center network all devices must support the same features. DCBX has been developed to assist in assuring that all the datacenter switching elements are in synchronous configuration with the same features enabled as appropriate.
- Enhanced Transmission Selection (ETS) 802.1qaz - ETS defines the queuing capability that a Data Center switch can support. This will ensure that the administrator can segment bandwidth between loss-less traffic and other types of traffic. ETS is a scheduling algorithm with customizable mapping of traffic classes to

queues.

- DCB Profiles - PFC, ETS, and DCBx configuration applied to switch ports through DCB profiles (concept similar to QSet Profiles that apply the QoS queue management configuration to switch ports). DCB profiles are based on the 802.1Q-REV/D1-5 standard to define how the switch classifies different traffic types and priority mappings and then groups those types into traffic classes.
 - Profiles also specify the Traffic Class Flow (TCF), which is LL (lossless; PFC initiated upstream) or nL ((lossy; PFC not initiated upstream).
 - DCB and QSet profiles are mutually exclusive on a port. Switch with DCB license comes up with DCB profile 8 on ports by default; DCBx and “willing” bit are enabled by default.
 - There are 10 predefined DCB profiles which cannot be modified or deleted. A total of 128 DCB profiles (1-10 predefined and 11-128 custom) are supported.

Edge Virtual Bridging (EVB)

Alcatel-Lucent’s Edge Virtual Bridging (EVB) follows the 802.1Qbg IEEE Standard. EVB is configured on the edge of a Data Center network and is used to automate the discovery of Virtual Machines and connect them to the correct network domains, either the backbone bridge service network, or the VLAN bridged network.

Note: In the current release EVB supports a single Edge Relay (ER) per port. The ER can operate as Virtual Edge Port Aggregator (VEPA) or Virtual Edge Bridge (VEB). EVB is supported under VLAN mode only.

Virtual Network Profiles/User Network Profiles

The Universal Network Profile (UNP) feature provides administrators with the ability to define and apply network access control to specific types of devices by grouping such devices according to specific matching profile criteria. This allows network administrators to create virtual machine network profiles (VNPs) and profiles for user devices from a unified framework of operation and administration. UNP is not limited to creating profiles to classify only certain types of devices. However, the following classification methods implemented through UNP functionality and profile criteria provide the ability to tailor profiles for specific devices (physical or virtual):

- MAC-based authentication using a RADIUS-capable server.
- Switch-wide classification rules to classify on source MAC or IP address (no authentication required).
- VLAN tag classification to create VLAN port or Service Access Point (SAP) associations based on the VLAN ID contained in device packets.
- Default UNP classification for untagged traffic or traffic not classified through other methods.

SAP / SPB-M Services - UNP provides a method for dynamically assigning network devices to VLAN domains or to a Shortest Path Bridging (SPB) service domain. A profile consists of configurable attributes. When a device sends traffic that matches such attributes, the device is then assigned to a VLAN or SPB service associated with the UNP. The UNP may also specify a QoS/ACL policy list that is subsequently applied to device traffic associated with the UNP VLAN or SPB service.

Customer Domains (Multi-tenancy) - Dynamic assignment of devices using UNP is achieved through port-based functionality that provides the ability to authenticate and classify device traffic. Authentication verifies the device identity and provides a UNP name. In the event authentication is not available or is unsuccessful, classification rules associated with the UNPs are applied to the traffic to determine the UNP VLAN or SPB service assignment.

Dynamic SAP Configuration - When device traffic is assigned to a service profile, UNP first checks the switch configuration to see if a SAP already exists for the VLAN tag, I-SID, and BVLAN ID values as defined in the service profile. If a SAP already exists with these values, the device traffic is classified into that SAP. If the SAP does not exist, the switch dynamically creates one for the profile traffic. In addition:

- If the I-SID value specified in the service profile does not exist, the switch will dynamically create the expected service and then the SAP as needed. This is similar to functionality in the VLAN domain, where UNP will dynamically create a VLAN based on the 802.1Q-tag of the device traffic.
- If the BVLAN ID specified in the service profile does not exist, the dynamic SAP is not created.

UNP also provides dynamic SAP configuration for UNP access port traffic that is not classified by a UNP service profile. Because there is no service profile associated with this traffic, UNP derives the necessary SAP parameters using pre-defined calculations. Refer to the Data Center Switching Guide for additional information.

Dynamically creating services and related SAPs is subject to available switch resources. If an attempt to dynamically create a service or SAP fails, for any reason, the MAC addresses classified for the service profile are learned as filtering.

NOTE: Dynamically created SAPs are not saved to the switch configuration file.

Release 7.3.1 adds support for UNP over MC-LAG on OS10K.

Layer 2 Feature Descriptions

Ethernet Ring Protection Verson 2 (ERPV2)

Ethernet Ring Protection (ERP) switching is a self-configuring algorithm that maintains a loop-free topology while providing data path redundancy and network scalability. ERP provides fast recovery times for Ethernet ring topologies by utilizing traditional Ethernet MAC and bridge functions.

This implementation of ERP is based on ITU-T G.8032 Version 2 and uses the ring Automatic Protection Switching (APS) protocol to coordinate the prevention of network loops within a bridged Ethernet ring. Loop prevention is achieved by allowing the traffic to flow on all but one of the links within the protected Ethernet ring. This link is blocked and is referred to as the Ring Protection Link (RPL). When a ring failure condition occurs, the RPL is unblocked to allow the flow of traffic to continue through the ring.

ERPV2 is enhanced to support multi rings and ladder networks. This introduces the concept of interconnection nodes, interconnected shared links, master rings and sub-rings. A shared link can only be part of one ring (i.e. the master ring). The sub-rings connected to the interconnection nodes are not closed and cannot use the shared links.

With the introduction of Version 2 the OmniSwitch supports the following:

- Backward compatibility with ERP v1
- Multi-ring and Ladder Networks
- Interconnection nodes

- Interconnected shared links
- Master Rings
- Sub-Rings
- Revertive Non-Revertive Mode

Not supported in 7.3.1:

- Multiple ERP instances per physical ring
- Administrative Commands – Forced Switch (FS), Manual Switch (MS), Clear for MS/FS
- Dual End Blocking

Layer 3 Feature Descriptions

VRF Management

This feature allows management services to be enabled or disabled in a VRF other than the default VRF. This allows for the creation of a single management VRF instance, a VRF per management service, or multiple VRFs for a service. Depending on the type of service there are different levels of management allowed as described below.

- Level 0 - The management service may only appear in the Default VRF.
- Level 1 - User may specify a single VRF that all management services can be configured in. For example, both RADIUS and LDAP can use vrf-1.
- Level 2 - Each management service or multiple management services can be configured for a different VRF. For example, RADIUS in vrf-1, LDAP in vrf-2, SNMP in vrf-3.
- Level 3 - A management service may appear in multiple VRFs. For example, SSH and Telnet in vrf-1 and vrf-2.

Level	Description	Telnet/SSH/SFTP/SCP	Radius/SNMP/HTTP/HTTPS/ NTP/LDAP/TACACS+/Syslog
0	Default VRF only	Yes	Yes
1	Single VRF for all services	Yes	Yes
2	Single VRF per service, each service can be on a different VRF	Yes	Yes
3	Multiple VRFs per service, any service on any VRF	Yes	No

VRF Route Leak

VRF route leaking provides the ability for devices/routers in one VRF to communicate with other VRFs in a controlled manner, without the need for any external devices. To achieve this the OmniSwitch now supports InterVRF routing by exporting routes to a Global Route Table (GRT) and then importing those routes into a separate VRF. In order to control the routes that are leaked the existing infrastructure of route-maps is used. VRF route leak supports the following:

Exporting Supports	Importing Supports
Match ip-address	Match ip-address
Match ip-nexthop	Match ip-nexthop
Match tag	Match tag
Match ipv4-interface	Match ipv4-interface
Match route-type	Match route-type
Set tag	
Set metric	

- Maximum of 128 routes in the GRT
- One route-map is allowed per VRF for export filtering
- One route-map is allowed for import filtering from each unique export VRF
- Route leaking supported on IPv4, IPv6 is not supported.
- Nesting is not supported.

Management Feature Descriptions

Virtual Chassis

Virtual Chassis is a group of chassis managed through a single management IP address. It provides both node level and link level redundancy for layer 2 and layer 3 services and protocols acting as a single device. The use of Virtual Chassis provides node level redundancy without the need to use redundancy protocols such as STP and VRRP between the edge and the aggregation/core layer.

The following are some key points regarding Virtual Chassis configuration:

- With the introduction of Virtual Chassis a switch can now run in two modes; Virtual Chassis or Standalone.
- A Virtual Chassis will consist of one master and one or more slave chassis. The election of a Master chassis can automatically be determined based on various chassis attributes.
- Running in Virtual Chassis mode will cause a change to the CLI requiring a chassis identifier to be used and displayed for the *slot/port* commands. (i.e. chassis/slot/port)
- A virtual chassis provides a single management IP address for a group of chassis that are acting as a single bridge or router.
- A Virtual Chassis can leverage an ISSU upgrade to help minimize network impact.
- The switches in the Virtual Chassis are created by inter-connecting them via standard single or aggregated 10G or 40G ports. 10G and 40G ports cannot be mixed within the same VFL.
- Each chassis participating in the Virtual Chassis must have a valid Advanced license to join the VC.

Virtual Chassis EMP IP Addresses

A Virtual Chassis has the following EMP IP addresses associated with it. These IP addresses can be used to access the entire Virtual Chassis or a specific chassis.

- **Virtual Chassis EMP Address** – The Virtual Chassis management IP address (EMP-VC)
- **Virtual Chassis Local EMP Address** - The local chassis management IP address (EMP-CHAS1 or EMP-CHAS2).

AAA/UNP Interaction with the EMP-VC

- Any authentication done with LDAP will not use the EMP-VC address.
- Any authentication done RADIUS or TACACS will use the EMP-VC address if configured.

Split chassis detection

The goal of split-chassis detection is to provide information in a virtual chassis environment which can be used to determine the state of other parts of a Virtual Chassis when a VFL has failed. A split chassis can occur when the VFL connection is broken but each of the chassis remains operational. To help prevent this scenario each chassis in the Virtual Chassis periodically sends information via the EMP port. All of the chassis participating in a Virtual Chassis should be able to communicate via the EMP port. When a VFL goes down, each chassis can still communicate with the other chassis to help resolve the split chassis scenario. If all VFL links go down then this protocol will detect and shutdown all user ports on the Slave chassis to avoid a loop. The user ports will automatically come up when the VFL connectivity is reestablished.

NOTE: If more than one Virtual Chassis is part of the same EMP out-of-band management network then each Virtual Chassis MUST have a unique Chassis-Group ID. Otherwise the split chassis protocol cannot differentiate between the two Virtual Chassis.

PFC Across the VFL

To support (PFC) across the VFL links of Virtual Chassis the links have to be configured in a certain manner. Only one VF Link is allowed from each port group based on the type of module being used to create the VFL. See the table below for the port group requirements:

Module	Port Grouping
OS10K-XNI-U32E	Port Group 1 (Ports 1-16) Port Group 2 (Ports 17-32)
OS10K-XNI-U16E/U16L	Port Group 1 (Ports 1-16)
OS10K-QNI-U8E	Port Group 1 (Ports 1-4) Port Group 2 (Ports 5-8)
OS10K-QNI-U4E	Port Group 1 (Ports 1-4)
Note: The OS6900 platform and the OS10K-XNI-U32S module are not affected	

Ethernet OAM

Ethernet OAM (Operation, Administration, and Maintenance) provides service assurance over a converged Ethernet network. Ethernet OAM focuses on two main areas that are most in need by service providers and are rapidly evolving in the standards bodies: Service OAM and Link OAM. These two OAM protocols have unique objectives but are complementary to each other. Service OAM provides monitoring and troubleshooting of end-to-end Ethernet service instances, while Link OAM allows a provider to monitor and troubleshoot an individual Ethernet link. The end-to-end service management capability is the most important aspect of Ethernet OAM for service providers.

This implementation of Ethernet Service OAM supports both IEEE 802.1ag Version 8.1 and ITU-T Y.1731 for connectivity fault management. Performance monitoring is provided by ITU-T Y.1731 using both oneway and two-way ETH-DM. Additionally, this implementation can perform delay measurement for both ITU-T Y.1731 and IEEE 802.1ag maintenance endpoints. Although both standards are supported, the OmniSwitch implementation uses the 802.1ag terminology and hierarchy for Ethernet CFM configuration.

Service Assurance Agent (SAA)

With SAAs, users can verify service guarantees, increase network reliability by validating network performance and proactively identify network issues. SAA uses active monitoring to generate traffic in a continuous, reliable, and predictable manner, thus enabling the measurement of network performance and health.

SAA - Ethernet OAM

ETH-LB/DMM can be used to measure delay and jitter by sending frames with DM information to the peer MEP and receiving frames with DM information from the peer MEP.

SAA - IP ping

IP SAAs enhances the service level monitoring to become IP application-aware by measuring both end-to-end and at the IP layer. It allows performance measurement against any IP addresses in the network (example Switch, Server, PC).

L2 SAA

L2 SAAs enhance the service level monitoring by enabling performance measurement against any L2 address within the provider network. This can be used to test connectivity in an SPB environment using both a MAC address and/or an ISID.

License Management

Some features require a software license and are restricted only to a licensed user. Purchasing a license along with an authorization code from Alcatel-Lucent is required. The authorization code is then used to generate a license file. The features below require the associated license.

- Advanced License – Required to support SPB and Virtual Chassis
- Data Center License – Required for Data Center Bridging Protocols (PFC,ETS,DCBX) and EVB.
- U16L License – Required to upgrade OS10K-XNI-U16L to a 16x10G line card.

Note: Advanced and Data Center licensed features are not supported over MC-LAG. Additionally, a reboot is required for any licensed features to take affect.

SFP+ Line Diags & EPP

EPP can assist in connecting with SFF-8431 non-compliant or electrically deficient devices. EPP can be used on some links to enhance the receive signal sampling resolution management and help to improve the link integrity to the link partner. The following steps should be followed to determine if EPP should be enabled:

- Check the current link quality - Check the current link quality of the interface. The Link-Quality can be Good, Fair, or Poor.
- Diagnose any link quality issues - If the Link Quality is not 'Good'. Perform a few basic troubleshooting steps to determine if the issue is with the link partner and whether enabling EPP can help improve the quality.
- Enable EPP - If it's determined that the issue is with the link partner, enable EPP.

Only certain transceivers support enabling EPP. Additionally, depending on the revision of the OmniSwitch, there are port restrictions due to the power requirements of enabling EPP as shown in the table below.

Product	Revision	EPP Support
OS6900-X20	B11	No restriction
	B10 or less	Only 5 ports can have EPP enabled
OS6900-X40	B11	No restriction
	B10 or less	Only 5 ports in 1st group of 20 and 5 ports in 2nd group of 20
Expansion Board	Any	No restrictions
10-Gigabit Transceivers	Any	Supported
1/40-Gigabit Transceivers	Any	Not Supported

Limited/Future Availability Features

The following software features are being introduced with the 7.3.1.R01 release as limited or early availability features. Some CLI and feature functionality may be available, however, they have not gone through the complete Alcatel-Lucent qualification process. For additional information please contact the Product Line Manager.

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

Limited/Future Availability Features Summary Table

Feature	Platform	License
Layer 3 Feature Support		
IS-IS – (IPv4)	OS10K/6900	Base

Limited/Future Availability Features Descriptions

Layer 3 Feature Support

IS-IS – (IPv4)

Intermediate System-to-Intermediate System (IS-IS) is an International Organization for Standardization (ISO) dynamic routing specification. IS-IS is a shortest path first (SPF), or link state protocol. Also considered an interior gateway protocol (IGP), IS-IS distributes routing information between routers in a single Autonomous System (AS) in IP environments. IS-IS chooses the least-cost path as the best path. It is suitable for complex networks with a large number of routers by providing faster convergence where multiple flows to a single destination can be simultaneously forwarded through one or more interfaces.

Existing Hardware Support

Existing Hardware - AOS 7.1.1.R01

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

OmniSwitch 10K Chassis

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

8 Slots – Network Interface Modules

2 Slots – Chassis Management Modules (Integrated Management and Chassis Fabric Module)

2 Slots – Chassis Fabric Modules

2 Slots – Fan Trays (Two fan trays required)

4 Slots – Power Supplies

OS10K-CMM

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

OS10K-CFM

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

OS10K-GNI-C48E

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

OS10K-GNI-U48E

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

OS10K-XNI-U32S

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

OS10K-PS-25A

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

OS10K-PS-24D

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

OS10K-Fan-Tray

Contains 12 individual variable-speed fans per tray.

Existing Hardware - AOS 7.2.1.R01

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

OmniSwitch 6900-X20

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

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

OmniSwitch 6900-X40

10-Gigabit Ethernet fixed configuration chassis in a 1U form factor with 40 SFP+ ports, two optional module slots, redundant AC or DC power and front-to-rear cooling. The switch includes:

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

OS-XNI-U4

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

OS-XNI-U12

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

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

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

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

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

OS6900-FT-F FanTray

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

Existing Hardware - AOS 7.2.1.R02

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

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

OmniSwitch 6900 Rear-to-Front Cooling

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

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

OS-QNI-U3 Module

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

OS-HNI-U6 Module

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

QSFP-40G-SR Transceiver

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

QSFP-40G-C Transceiver

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

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

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

Note: This power supply is differentiated from the front-to-rear power supplies by purple coloring.

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

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

Note: This power supply is differentiated from the front-to-rear power supplies by purple coloring.

OS6900-FT-R FanTray

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

Note: This fan tray is differentiated from the front-to-rear fan tray by an **R->F** label and purple coloring.

Existing Software Support

Existing Software Features - AOS 7.1.1.R01

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

AOS 7.1.1. R01 Feature Summary Table

Feature	Platform	Software Package
Manageability Feature Support		
CLI	OS10K	Base
Ethernet Interfaces	OS10K	Base
ISSU	OS10K	Base
Multiple VRF Routing and Forwarding	OS10K	Base
Network Time Protocol (NTP)	OS10K	Base
Pause Control/Flow Control	OS10K	Base
Remote Access FTP SCP SSH/SFTP Telnet TFTP	OS10K	Base
Smart Continuous Switching Hot Swap Management Module Failover Power Monitoring Redundancy	OS10K	Base
SNMP	OS10K	Base
Software Rollback – Multi-Image/Multi-Config	OS10K	Base
Storm Control	OS10K	Base
Text File Configuration	OS10K	Base
UDLD	OS10K	Base
USB Support	OS10K	Base
Web-Based Management (WebView)	OS10K	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS10K	Base
802.1Q	OS10K	Base
Configurable Hash Mode	OS10K	Base
Link Aggregation –Static and LACP (802.3ad)	OS10K	Base
Multi-Chassis Link Aggregation	OS10K	Base
Source Learning	OS10K	Base

Feature	Platform	Software Package
Spanning Tree <ul style="list-style-type: none"> • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard 	OS10K	Base
VLANs	OS10K	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS10K	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS10K	Base
BGP4 with Graceful Restart	OS10K	Base
DNS Client	OS10K	Base
GRE	OS10K	Base
IP Multicast Routing	OS10K	Base
IP Multicast Switching (IGMP)	OS10K	Base
IP Multicast Switching (Proxying)	OS10K	Base
IP Multinetting	OS10K	Base
IP Route Map Redistribution	OS10K	Base
IP-IP Tunneling	OS10K	Base
OSPFv2	OS10K	Base
RIPv1/v2	OS10K	Base
Routing Protocol Preference	OS10K	Base
Server Load Balancing	OS10K	Base
VRRPv2	OS10K	Base
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

Feature	Platform	Software Package
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS10K	Base
Ingress and egress bandwidth shaping	OS10K	Base
Policy Based Routing	OS10K	Base
Tri-Color Marking	OS10K	Base
Multicast Feature Support		
DVMRP	OS10K	Base
IGMP Multicast Group Configuration Limit	OS10K	Base
IGMP Relay	OS10K	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS10K	Base
L2 Static Multicast Address	OS10K	Base
PIM / PIM-SSM (Source-Specific Multicast)	OS10K	Base
Monitoring/Troubleshooting Feature Support		
DDM - Digital Diagnostic Monitoring	OS10K	Base
Health Statistics	OS10K	Base
Ping and Traceroute	OS10K	Base
Policy Based Mirroring	OS10K	Base
Port Mirroring	OS10K	Base
Port Monitoring	OS10K	Base
Remote Port Mirroring	OS10K	Base
Rmon	OS10K	Base
sFlow	OS10K	Base
Switch Logging and Syslog	OS10K	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS10K	Base
Ethernet Services	OS10K	Base
L2 Control Protocol Tunneling (L2CP)	OS10K	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS10K	Base
Account & Password Policies	OS10K	Base
Admin User Remote Access Restriction Control	OS10K	Base
ARP Defense Optimization	OS10K	Base
ARP Poisoning Detect	OS10K	Base
Authenticated Switch Access	OS10K	Base
IP DoS Filtering	OS10K	Base
Learned Port Security (LPS)	OS10K	Base
Policy Server Management	OS10K	Base

Existing Software Features – AOS 7.2.1.R01

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

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

AOS 7.2.1.R01 Feature Summary Table

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

Feature	Platform	License
Spanning Tree <ul style="list-style-type: none"> • 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard 	OS6900	Base
Universal Network Profiles (UNP)	OS6900	Base
VLANs	OS6900	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS6900	Base
DHCP / UDP DHCP Relay/Option-82 Per-VLAN UDP Relay	OS6900	Base
BGP4 with Graceful Restart	OS6900	Advanced
DNS Client	OS6900	Base
GRE	OS6900	Base
IP Multicast Routing	OS6900	Advanced
IP Multicast Switching (IGMP)	OS6900	Base
IP Multicast Switching (Proxying)	OS6900	Base
IP Multinetting	OS6900	Base
IP Route Map Redistribution	OS6900	Base
IP-IP Tunneling	OS6900	Base
OSPFv2	OS6900	Advanced
RIPv1/v2	OS6900	Base
Routing Protocol Preference	OS6900	Base
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

Feature	Platform	License
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS6900	Base
Ingress and egress bandwidth shaping	OS6900	Base
Policy Based Routing	OS6900	Advanced
Tri-Color Marking	OS6900	Base
Multicast Feature Support		
DVMRP	OS6900	Advanced
IGMP Multicast Group Configuration Limit	OS6900	Base
IGMP Relay	OS6900	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS6900	Base
L2 Static Multicast Address	OS6900	Base
PIM / PIM-SSM (Source-Specific 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
IP DoS Filtering	OS6900	Base
Learned Port Security (LPS)	OS6900	Base
Policy Server Management	OS6900	Base

Existing Software Features – AOS 7.2.1.R02

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

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

Feature	Platform	License
Layer 2 Feature Support		
High Availability VLAN <ul style="list-style-type: none"> Added support for OS10K HA-VLAN with MCLAG 	OS10K OS10K/6900	Base Base
Multi-Chassis Link Aggregation <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> Added support for OS10K 	OS10K	Base
Universal Network Profiles <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> Added support for profiles 2-4 Added support for WRED 	OS10K/6900 OS6900	Base Base
Security		
Learned Port Security Enhancements	OS10K/6900	Base

Existing Software Features and Enhancements Descriptions as released in 7.2.1.R02

Layer 2 Feature Descriptions

High Availability VLAN (HA-VLAN)

This release enhances the HA-VLAN feature as described below:

OS10K Support

HA-VLAN is now supported on the OS10K.

MC-LAG Support

HA-VLAN is now supported on MC-LAG configuration on both the OS10K and 6900. See the MC-LAG section for feature interaction.

Multi-Chassis Link Aggregation

This release enhances the MC-LAG feature as described below:

Configurable Chassis Group ID (Multiple MC-LAG Domains)

By default, the Chassis Group ID is set to "0". In a network environment with more than one MC-LAG domain, such as a back-to-back topology, the administrator can now configure each domain with its own unique Group ID.

The group ID is used to generate a globally unique virtual MAC address for each multi-chassis domain to avoid duplicate MAC addresses in a network that may contain more than one MCLAG domain configuration.

Standalone port in VIP VLAN

Prior to this release only MC-LAG ports could be members of the VIP VLAN. This release allows any port to be a member of the VIP VLAN. In addition, assigning MCLAG ports to standard VLANs (non-VIP VLANs) is supported.

Interaction with Server Load Balancing

MCLAG now supports the use of Server Load Balancing (SLB) in a multi-chassis configuration.

The SLB configuration must be the same on both MCLAG peer switches. Any inconsistencies in the configuration between the two switches could impact the flow of traffic, especially in a failover scenario.

There is no synchronization of the SLB operation between the two peer switches. This means that as servers become reachable or unreachable, a period of time may occur during which the hashing is different on each peer switch.

Note: The SLB VIP cannot be the MCLAG VIP.

Interaction with High Availability VLANs (HA-VLANs)

MCLAG now supports the use of HA-VLANs in a multi-chassis configuration:

- The High Availability VLAN (HAVLAN) configuration must be the same on both MCLAG peer switches. Enhancements to the HAVLAN show commands provide information to determine the status of a HAVLAN consistency check between the two multi-chassis peer switches, the reason for any inconsistency detected, and if the VFL link is in use by the server cluster.

- When the HAVLAN configuration and MCLAG configuration are both up and running, the HAVLAN feature performs a consistency check to verify that the server cluster configuration is the same on both MCLAG peer switches.
- If the server cluster is not operationally up on both peer switches, then the server cluster configuration is only applicable to the local switch. In this case, an HAVLAN consistency check is not started.
- In the case of an L3 server cluster, the VLAN used for the IP interface must exist on both peer switches and must also be a VIP VLAN.
- An SNMP trap notification is sent when the HAVLAN consistency check detects a synchronization error between the two peer switches (server cluster configuration does not match on both peers).

Interaction with Universal Network Profiles

MCLAG now supports the use of User Network Profiles (UNP) in a multi-chassis configuration:

- The Universal Network Profile (UNP) configuration must be the same on both MCLAG peer switches. Any inconsistencies in the configuration between the two switches could effect the traffic flow, especially in a failover scenario. Enhancements to the MCLAG show commands provide information that indicates the status of a UNP consistency check between the two multi-chassis peer switches.
- When UNP detects that the MCLAG configuration is up and operational, UNP performs a consistency check to verify that the UNP configuration is the same on both multi-chassis peer switches. The following conditions will also trigger a consistency check:
 - If there is any change to the UNP configuration.
 - If the MCLAG operation goes down and comes back up.
- When the UNP configuration is up and synchronized between the peer switches, disabling and enabling all UNPs on MCLAG aggregates is recommended. Doing so will clear any transient MAC addresses from the configuration.
- There is no automatic correction of any UNP configuration inconsistency, it is up to the administrator to ensure the necessary changes are made to bring the configuration back into synch between the two MCLAG peer switches.

MVRP

MVRP is now supported on the OS10K.

Universal Network Profiles (UNP)

This release enhances the UNP feature as described below:

Interaction with MVRP

- UNP supports a dynamic profile configuration option. When this option is enabled, tagged packets received on UNP ports that are enabled to trust the VLAN tag, are classified based on the VLAN tag of the packet. If the VLAN tag matches a MVRP VLAN on the switch and the MVRP VLAN is not already assigned to a profile:
 - A new profile is automatically created and associated with the MVRP VLAN.
 - The MVRP VLAN is converted to a UNP dynamic VLAN if the UNP dynamic VLAN configuration option is also enabled for the switch.
- MVRP is not supported on UNP ports; however, both features can co-exist on the same switch. The recommended configuration is to have UNP dynamically create VLAN-port-associations on edge ports while MVRP propagates the dynamic VLANs down and up stream.

Interaction with Link Aggregation / MC-LAG

- UNP can now be configured on Link Aggregate ports.
- UNP with MC-LAG (See MC-LAG section).

Interaction with LPS

- UNP can now be configured on the same port with LPS. LPS classification is applied first, then UNP classification rules.

Layer 3 Feature Descriptions**ECMP**

The OmniSwitch now supports 16 ECMP routes for IPv6.

QoS Feature Descriptions**VFC/VoQ Profiles**

- Four pre-defined Qset Profiles (QSPs) are now supported.
- WRED is supported on an OS6900.
- WRED is not supported on the OS10K.
- The OS10K and OS6900 CLI syntax for configuring QSet profiles and instances is now the same.

Security Features Descriptions**Learned Port Security Enhancements**

The following Learned Port Security (LPS) enhancements have been added:

- Universal command for assigning static MAC addresses regardless of port state.
- LPS now continues to learn filtering MAC addresses after the learning window has expired, but only up to the configured filtering MAC address limit.
- A new type of static MAC address (pseudo-static) is maintained. A pseudo-static MAC address is not user-configured; it is a dynamically learned MAC address that is treated the same as a regular static address (will not age out even if the learning window closes). However, the pseudo-static MAC is not saved in the running configuration.
- New parameter options for the LPS **port-security shutdown** CLI command.
 1. **No Aging of Learned MAC Addresses.** A new **no-aging** parameter specifies whether or not LPS will learn MAC addresses as “pseudo-static” addresses.
 2. **Convert MAC Addresses to Static MACs.** A new **convert-to-static** parameter specifies whether or not pseudo-static and dynamically learned MAC addresses are converted to static MAC addresses when the learning window time expires.
 3. **Learning Window Start at Boot-up.** A new **boot-up** parameter specifies whether or not LPS will start the learning window time when the switch boots up.
- New **admin-state** parameter for the **port-security** CLI command. This parameter is used to enable, disable, or lock an LPS port. In addition, the **port-security** command will now accept a range of ports.

- New **brief** parameter for the **show port-security** CLI command. This parameter is used to provide a summary of the LPS status, configuration, and MACs learned on all the LPS ports.
- The VLAN ID bound to an LPS static MAC address is automatically updated when the default VLAN for the LPS port is changed.
- Duplicate LPS static MAC addresses are now allowed on different ports within the same VLAN. However, dynamic MAC addresses that match a configured static MAC address within the same VLAN are not learned.
- The “Bridge MAC Learned” and “LPS Violation” SNMP traps now have three fields of information: port number, VLAN ID, and MAC address.
- A new LPS shutdown violation mode, “discard”, is now supported. This mode administratively disables the port, but the port remains physically up. The “shutdown” and “restricted” modes are still supported.

NOTE: The convert to static option is not supported on a port that also has UNP enabled since UNP and static MACs cannot be configured on the same port.

Existing Software Feature Descriptions

Manageability Feature Support

Command Line Interface (CLI)

The command line interface (CLI) is a text-based configuration interface that allows configuration of switch applications and to view switch statistics. Each CLI command applicable to the switch is defined in the CLI Reference guide. All command descriptions listed in the Reference Guide include command syntax definitions, defaults, usage guidelines, example screen output, and release history. The CLI uses single-line text commands that are similar to other industry standard switch interfaces.

Ethernet Interfaces

The OmniSwitch supports Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet ports. This includes configuration of basic line parameters, gathering of statistics and responding to administrative enable/disable requests. Configurable parameters include: autonegotiation, trap port link messages, flood control, line speed, duplex mode, resetting statistics counters, and maximum and peak flood rates.

In-Service Software Upgrade (ISSU)

The In-Service Software Upgrade (ISSU) feature is used to upgrade the CMM and NI images running on an OmniSwitch 10K with minimal disruption to data traffic. The images can be upgraded on a fully synchronized, certified, and redundant system running an ISSU capable build without requiring a reboot of the switch.

- ISSU is not supported using WebView in this release.
- ISSU upgrades are only supported within the same software branch. For example, if a switch is running 7.1.1.###.R01 then only 7.1.1.###.R01 images can be used to perform an ISSU patch. If a switch is running 7.1.1.###.R01 then ISSU is not supported with 7.2.1.###.R02 images.

License Management

Some features require a software license and are restricted only to a licensed user. Purchasing a license along with an authorization code from Alcatel-Lucent is required. The authorization code is then used to generate a license file. The features below require an **Advanced** license.

Layer 3		Multicast	Other
OSPF v2/v3	VRRP	DVMRP	IPSec IPv6
BGP	VRRP v3	PIM	VRF
MP-BGP	RIPng	PIM-SM IPv6	
Policy Based Routing			

Advanced License Features

Multiple Virtual Routing and Forwarding (Multiple-VRF)

The Multiple Virtual Routing and Forwarding (VRF) feature provides the ability to configure separate routing instances on the same switch. Similar to using VLANs to segment Layer 2 traffic, VRF instances are used to segment Layer 3 traffic. Some of the benefits of using the Multiple VRF feature include the following:

Multiple routing instances within the same physical switch. Each VRF instance is associated with a set of IP interfaces and creates and maintains independent routing tables. Traffic between IP interfaces is only routed and forwarded within those interfaces/routes that belong to the same VRF instance.

Multiple instances of IP routing protocols, such as static, RIP, IPv4, BGPv4, and OSPFv2 on the same physical switch. An instance of each type of protocol operates within its own VRF instance.

The ability to use duplicate IP addresses across VRF instances. Each VRF instance maintains its own IP address space to avoid any conflict with the service provider network or other customer networks.

Separate IP routing domains for customer networks. VRF instances configured on the Provider Edge (PE) are used to isolate and carry customer traffic through the shared provider network.

The Multiple VRF feature uses a context-based command line interface (CLI). When the switch boots up, a default VRF instance is automatically created and active. Any commands subsequently entered apply to this default instance. If a different VRF instance is selected, then all subsequent commands apply to that instance. The CLI command prompt indicates which instance is the active VRF CLI context by adding the name of the VRF instance as a prefix to the command prompt (for example, **vrf1: ->**).

VRF - QoS

Allows QoS policy configuration by adding a field in the policy condition to allow a VRF instance to be specified. The VRF classification can be combined with any existing condition and allows for the configuration of VRF aware policy rules.

VRF - Switch Authentication

This feature allows a RADIUS server to be placed in a VRF other than the default VRF. This allows for the creation of a Management VRF instance where all authentication servers can be placed. Authentication servers may also be left in the non-default VRF instance.

VRF - Switch Access and Utilities

Telnet and SSH are VRF aware. This feature applies only to outgoing Telnet and SSH connections from any VRF instance, incoming requests always go to the default VRF instance. Additionally, the ping and traceroute utilities are also VRF aware.

VRF - VRRP

Allows for the configuration of independent VRRP instances in multiple VRFs. The existing VRRP commands and syntaxes (including show commands and outputs) are now accessible in a “VRF” context. VRRP instances can be configured independently of one another on as many VRFs as the underlying platform supports. Each VRRP/VRF instance receives, sends, and processes VRRP packets independently of VRRP instances running in other VRFs.

VRF - UDP/DHCP Relay

VRF support for UDP/DHCP Relay allows for the configuration and management of relay agents and servers within the context of a VRF instance. However, the level of VRF support and functionality for individual UDP/DHCP Relay commands falls into one of the following three categories:

VRF-Aware commands. These commands are allowed in any of the VRF instances configured in the switch.

The settings in one VRF are independent of the settings in another VRF. Command parameters are visible and configurable within the context of any VRF.

Global commands. These commands are supported only in the default VRF, but are visible and applied to all VRF instances configured in the switch. This command behavior is similar to how command parameters are applied in the per-VLAN DHCP Relay mode. For example, the maximum hops value configured in the default VRF is applied to all DHCP Relay agents across all VRF instances. This value is not configurable in any other VRF instance.

Default VRF commands. These commands are supported only in the default VRF and are not applied to any other VRF instance configured in the switch. For example, per-VLAN mode and boot-up commands fall into this category.

Note: A switch running multiple VRF instances can only be managed with SNMPv3. A context must be specified that matches the VRF instance to be managed.

VRF – PIM and DVMRP

PIM-DM, PIM-SM, and DVMRP are VRF aware.

Network Time Protocol (NTP)

The Network Time Protocol (NTP) is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver. The OmniSwitch NTP client/server software will be able to respond to NTP client requests, and establish a client/server peering relationship.

NTP Loopback

Provides ability to configure a permanent source IP interface to be used when sending packets. The source IP interface can be the Loopback0 address or an existing IP interface on the switch.

Pause Control/Flow Control

PAUSE frames are used to pause the flow of traffic between two connected devices when traffic congestion occurs. PAUSE frame flow control provides the ability to configure whether or not the switch will transmit and/or honor PAUSE frames on an active interface. This feature is only supported on interfaces configured to run in full-duplex mode.

In addition to configured PAUSE frame flow control settings, this feature also works in conjunction with auto-negotiation to determine operational transmit/receive settings for PAUSE frames between two switches. Note that the configured PAUSE frame flow control settings are overridden by the values that are determined through auto-negotiation.

Remote Access

- **File Transfer Protocol (FTP)**

FTP can be used to transfer files to and from an OmniSwitch. The OmniSwitch can act as either a FTP client or server.

- **Secure Copy (SCP)**

SCP is used in a secure manner between hosts on the network. The scp utility performs encrypted data transfers using the Secure Shell (SSH) protocol. In addition, scp uses available SSH authentication and security features, such as prompting for a password if one is required.

- **Secure Shell (SSH)/Secure FTP (SFTP)**

The Secure Shell feature provides a secure mechanism that allows you to log in to a remote switch, to execute commands on a remote device, and to move files from one device to another. Secure Shell provides secure, encrypted communications even when your transmission is between two untrusted hosts or over an unsecured network.

The OmniSwitch includes both client and server components of the Secure Shell interface and the Secure Shell FTP file transfer protocol. SFTP is a subsystem of the Secure Shell protocol. All Secure Shell FTP data are encrypted through a Secure Shell channel.

- **Telnet**

Telnet can be used to log into the switch from a remote station. The OmniSwitch can act as either a Telnet client or server.

- **Trivial File Transfer Protocol (TFTP)**

TFTP, a client-server protocol, can be used to transfer files between a TFTP server and client. TFTP client functionality on the OmniSwitch is used to download files from or upload files to the TFTP server.

Hardware Resiliency

All OmniSwitch models support N+1 redundant, hot-swappable AC and DC power supplies. The primary and backup power supply units are internal, but removable allowing for easier maintenance and replacement. There is no interruption of service when a new power supply is installed or an old one replaced. Additionally the switch supports hot-swapping of the fan trays, plug-in modules and NIs of the same type.

SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that allows communication between SNMP managers and SNMP agents on an IP network. Network administrators use SNMP to monitor network performance and to solve network problems. SNMP provides an industry standard communications model used by network administrators to manage and monitor their network devices. The OmniSwitch supports SNMPv1, SNMPv2, and SNMPv3.

Software Rollback – Multi-image/Multi-Config

The directory structure inherent in an OmniSwitch switch allows for a switch to return to a previous, more reliable version of image or configuration files.

Changes made to the configuration file may alter switch functionality. These changes are not saved unless explicitly done so by the user. If the switch reboots before the configuration file is saved, changes made to the configuration file prior to the reboot are lost.

Likewise, new image files should be placed in a non-certified directory first. New image or configuration files can be tested to decide whether they are reliable. Should the configuration or image files prove to be less reliable than their older counterparts in the *certified* directory, then the switch can be rebooted from the *certified* directory, and “rolled back” to an earlier version.

Once the contents of the working directory are established as good files, then these files can be saved to the *certified* directory and used as the most reliable software to which the switch can be rolled back to in an emergency situation.

- **Multi-Image/Multi-Config**

The Multi-Image/Multi-Config feature allows for multiple switch configurations to be saved to user-defined directories. These configurations can be used to store additional switch configurations that can be loaded at any time.

Storm Control

The OmniSwitch storm/flood control feature for broadcast, multicast, and unknown unicast traffic can be limited based on bits-per-second, percentage of the port speed, or packets per second.

Text File Configuration

The text file configuration feature allows you to configure the switch using an ASCII-based text file. You may type CLI commands directly into a text document to create a configuration file. This file resides in the switch's file system. You can create configuration files in the following ways:

- You may create, edit and view a file using a standard text editor (such as Microsoft NotePad) on a workstation. The resulting configuration file is then uploaded to the switch.
- You can invoke the switch's CLI snapshot command to capture the switch's current configuration into a text file.

- You can use the switch's text editor to create or make changes to a configuration file.

UDLD - Fiber and Copper

The unidirectional link detection protocol is a protocol that can be used to detect and disable malfunctioning unidirectional Ethernet fiber or copper links. Errors due to improper installation of fiber strands, interface malfunctions, media converter faults, etc can be detected and the link can be disabled. It operates at Layer 2 in conjunction with IEEE 802.3's existing Layer 1 fault detection mechanisms.

USB Support

The USB port can be used with an Alcatel-Lucent certified USB Flash drive (OS-USB-FLASHDR) to provide the following functions:

- Disaster Recovery – The switch can boot from the USB drive if it is unable to load AOS from flash.
- Upload / Download Image and Configuration Files - To create or restore backup files.
- Upgrade Code - Upgrade code with the image files stored on the USB drive.

Web-Based Management (WebView)

The switch can be monitored and configured using WebView, Alcatel-Lucent's web-based device management tool. The WebView application is embedded in the switch and is accessible via the following web browsers:

IE6, IE7, IE8 for Windows XP

IE8, Firefox Mozilla 3.6 on Vista

WebView contains modules for configuring all features in the switch. Configuration and monitoring pages include context-sensitive on-line help.

Layer 2 Feature Support

802.1AB MED Extensions

The Link Layer Discovery Protocol-Media Endpoint Discover (LLDP-MED) is designed to extend IEEE 802.1AB functionality to exchange information such as VLANs and power capabilities. 802.1AB MED adds support for Network Policy and Inventory Management.

802.1Q

802.1Q is an IEEE standard for sending frames through the network tagged with VLAN identification. 802.1Q tagging is the IEEE version of VLANs. It is a method of segregating areas of a network into distinct VLANs. By attaching a label, or tag, to a packet, it can be identified as being from a specific VLAN or identified as being destined for a specific VLAN.

Configurable Hash Mode

Hashing helps in achieving better load balancing on the switch for features such as Link Aggregation, ECMP and Server Load Balancing. Depending on the OmniSwitch configuration, this feature allows the hashing mode to be configured to help improve switch load balancing performance.

There are two hashing algorithms available, Brief Mode or Extended Mode. In brief mode UDP/TCP ports will not be included in the hashing algorithm and only source IP and destination IP addresses are considered. Extended mode allows for additional bits to be used in the hashing algorithm as well as providing the option of allowing UDP/TCP ports to be included in the hashing algorithm resulting in more efficient load balancing.

Default Hashing Mode and Recommendations

Platform	Default Hashing Mode
OS6900	Brief
OS10K	Extended

- Changing the hash mode affects all features that rely on hashing, including Link Aggregation, ECMP and Server Load Balancing. Changing the hash mode per feature is not supported.
- Server Load Balancing uses dynamic port assignment, therefore it is not recommended to enable the TCP/UDP port hashing option with extended mode when SLB is configured on the switch.

High Availability -VLAN

High availability (HA) VLANs send traffic intended for a single destination MAC address to multiple switch ports. The HA VLAN feature on the OmniSwitch provides a flexible way to connect server cluster nodes directly to the ingress network. This involves multicasting the service requests on the configured ports. The multicast criteria is configurable based on destination MAC and destination IP address. Egress ports can be statically configured on a server cluster or they can be registered by IGMP reports. The server cluster feature on the OmniSwitch multicast the incoming packets based on the server cluster configuration on the ports associated with the server cluster.

HA VLAN Operational Modes

There are two modes of implementation of server clusters using HA VLANs.

Layer 2 - The server cluster is attached to a L2 switch on which the frames destined to the cluster MAC address are flooded on all interfaces by configuring static MAC addresses.

Layer 3 - The server cluster is attached to a L3 switch on which the frames destined to the server cluster IP address are routed to the server cluster IP and then flooded on all interfaces by configuring static ARP entries.

Link Aggregation - Static & LACP (802.3ad)

Alcatel-Lucent's link aggregation software allows you to combine several physical links into one large virtual link known as a link aggregation group. Using link aggregation can provide the following benefits:

- **Scalability.** You can configure up to 128 link aggregation groups.
- **Reliability.** If one of the physical links in a link aggregate group goes down, the link aggregate group can still operate.
- **Ease of Migration.** Link aggregation can ease the transition from a Gigabit Ethernet backbone to a 10-Gigabit Ethernet backbone.
- **Interoperability with Legacy Switches.** Static link aggregation can interoperate with OmniChannel on legacy switches.

Non-Unicast Load Balancing on Link Aggregation

The OmniSwitch supports load balancing of non-unicast (broadcast, multicast, flood) traffic over Link Aggregation. Hashing criteria is configurable. By default the hashing keys are derived from the flow-based attributes listed below:

Uses source and destination IP addresses for IP frames.

Uses source and destination MAC address for non-IP frames.

Multi-Chassis Link Aggregation

The Multi-Chassis Link Aggregation feature (MC-LAG) provides resiliency at the edge of the network by enabling dual homing of any standards-based edge switches to a pair of aggregation switches to provide a Layer 2 multipath infrastructure. The feature allows links that are physically connected to two different OmniSwitches to appear as a single link aggregation group to a third edge device. MC-LAG enables a device to form a logical link aggregation (LAG) interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multihoming support, and loop-free Layer 2 network without running Spanning Tree Protocol (STP).

Note: MC-LAG between an OS6900 and OS10K is not supported.

MVRP - Multiple VLAN Registration Protocol

Multiple VLAN Registration Protocol as defined in IEEE 802.1ak is intended as a replacement to GVRP by offering more scalable capabilities for large bridged networks. MVRP's general operation is similar to GVRP in that it controls and signals dynamic VLAN registration entries across the bridged network. MVRP addresses these major areas for improvements over GVRP:

- Improved PDU format to fit all 4094 VLANs in a single PDU.
- Reduced unnecessary flushing from STP topology changes that do not impact the Dynamic VLAN topology

Source Learning

Source Learning builds and maintains the MAC address table on each switch. New MAC address table entries are created in one of two ways: they are dynamically learned or statically assigned. Dynamically learned MAC addresses are those that are obtained by the switch when source learning examines data packets and records the source address and the port and VLAN it was learned on. Static MAC addresses are user defined addresses that are statically assigned to a port and VLAN.

In addition, Source Learning also tracks MAC address age and removes addresses from the MAC address table that have aged beyond the configurable aging timer value.

Accessing MAC Address Table entries is useful for managing traffic flow and troubleshooting network device connectivity problems.

- **Disable Learning on a Per Port Basis**
Provides the option to disable source learning on a per port basis. This feature is only supported on “hardware learning” ports and is not supported on mobile ports, LPS ports or Access Guardian ports. The feature is also supported for Link Aggregation where all ports in the aggregate are set to disable source learning. Configuration of static mac-addresses on such ports is still allowed.
- **Disable MAC Learning on a Per VLAN Basis**
Provides the option to disable source learning for all the ports of a VLAN. This feature is meant to be used on a ring topology where a VLAN only contains two ports. It is recommended to have only 2 ports in a VLAN that has source learning disabled.

Note: Disabling of source learning for a VLAN is not supported on the OS6900.

Spanning Tree

The OmniSwitch provides support for the 802.1w Rapid Spanning Tree Algorithm and Protocol (RSTP) and the 802.1D Spanning Tree Algorithm and Protocol (STP). Spanning Tree protocols ensure that there is always only one data path between any two switches for a given Spanning Tree instance to prevent network loops.

MSTP) is only available when the flat mode is active for the switch. The flat mode applies a single spanning tree instance across all VLAN port connections on a switch. MSTP allows the configuration of Multiple Spanning Tree Instances (MSTIs) in addition to the CST instance. Each MSTI is mapped to a set of VLANs. As a result, flat mode can now support the forwarding of VLAN traffic over separate data paths.

802.1D STP and 802.1w RSTP are available in both the flat and 1x1 mode. However, when using 802.1D or 802.1w in the flat mode, the single spanning tree instance per switch algorithm applies. Note that 802.1w is now the default Spanning Tree protocol for the switch regardless of which mode is active. This default value will apply to future releases as well.

Multiple Spanning Tree Protocol (MSTP)

Multiple Spanning Tree Protocol (MSTP) is a combination of the 802.1D 2004 and 802.1S protocols. This implementation of Q2005 also includes improvements to edge port configuration and provides administrative control to restrict port role assignment and the propagation of topology change information through bridge ports.

PVST+ Interoperability

The current Alcatel-Lucent 1x1 Spanning Tree mode has been extended to allow all user ports on an OmniSwitch to transmit and receive either the standard IEEE BPDUs or proprietary PVST+ BPDUs. An OmniSwitch can have ports running in either 1x1 mode when connecting to another OmniSwitch, or PVST+ mode simultaneously.

- It is mandatory that all the Cisco switches have the Mac Reduction Mode feature enabled.
- Priority values can only be assigned in multiples of 4096 to be compatible with the Cisco MAC Reduction mode.
- In a mixed OmniSwitch and Cisco environment, it is highly recommended to enable PVST+ mode on all OmniSwitches in order to maintain the same root bridge for the topology.
- Alcatel-Lucent's PVST+ interoperability mode is not compatible with a switch running in PVST mode.
- The same default path cost mode, long or short, must be configured the same way on all switches.

Universal Network Profile (UNP)

A Universal Network Profile (UNP) defines network access controls and resources for one or more physical or logical devices. Each device that is assigned to a specific profile is granted network access based on the profile criteria, instead of on an individual MAC address, IP address, or port.

Assigning devices to a profile provides greater flexibility and scalability across the network. Administrators can use profiles to group devices according to function. All devices assigned to the same UNP become members of that profile group. The UNP then determines what network access controls and resources are available to a group of devices, regardless of source subnet, VLAN or other characteristics.

A UNP consists of the following attributes:

UNP Name. The UNP name is obtained from the RADIUS server and mapped to the same profile name configured on the switch. The switch profile then identifies three attribute values: VLAN ID, Host Integrity Check (HIC) status, and a QoS policy list name.

VLAN ID. All members of the profile group are assigned to the VLAN ID specified by the profile.

QoS Policy List Name. Specifies the name of an existing list of QoS policy rules. The rules within the list are applied to all members of the profile group to enforce access to network resources. Only one policy list is allowed per profile, but multiple profiles may use the same policy list.

VLANs

One of the main benefits of using VLANs to segment network traffic, is that VLAN configuration and port assignment is handled through switch software. This eliminates the need to physically change a network device connection or location when adding or removing devices from the VLAN broadcast domain. The VLAN management software handles the following VLAN configuration tasks:

- Creating or modifying VLANs.
- Assigning or changing default VLAN port associations (VPAs).
- Enabling or disabling VLAN participation in the current Spanning Tree algorithm.
- Enabling or disabling classification of mobile port traffic by 802.1Q tagged VLAN ID.
- Enabling or disabling VLAN authentication.
- Enabling or disabling unique MAC address assignments for each router VLAN defined.
- Displaying VLAN configuration information.

IPv4 Feature Support

Internet Protocol (IP) is a network-layer (Layer 3) protocol that contains addressing and control information that allow packets to be forwarded on a network. IP is the primary network-layer protocol in the Internet protocol suite. Along with the Transmission Control Protocol (TCP), IP represents the heart of the Internet protocols. IP is associated with several Layer 3 and Layer 4 protocols. These protocols are built into the base code loaded on the switch and they include:

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Bootstrap Protocol (BOOTP)/Dynamic Host Configuration Protocol (DHCP)
- Simple Network Management Protocol (SNMP)
- Remote Access
- Address Resolution Protocol (ARP)
- Internet Control Message Protocol (ICMP)
- RIP I / RIP II
- OSPF
- BGP
- Static Routes

The base IP software allows one to configure an IP router interface, static routes, a default route, the Address Resolution Protocol (ARP), the router primary address, the router ID, the Time-to-Live (TTL) Value, IP-directed broadcasts, and the Internet Control Message Protocol (ICMP). In addition, this software allows one to trace an IP route, display Transmission Control Protocol (TCP) information, and display User Datagram Protocol (UDP) information.

Bi-Directional Forwarding Detection (BFD)

Bidirectional Forwarding Detection (BFD) is a hello protocol that can be configured to interact with routing protocols for the detection of path failures and can reduce the convergence time in a network. BFD is supported with the following Layer 3 protocols: BGP, OSPF, VRRP Tracking and Static Routes.

When BFD is configured and enabled, BFD sessions are created and timers are negotiated between BFD neighbors. If a system does not receive a BFD control packet within the negotiated time interval, the neighbor system is considered down. Rapid failure detection notices are then sent to the routing protocol, which initiates a routing protocol recalculation. This process can reduce the time of convergence in a network.

BGP4

The Border Gateway Protocol (BGP) is an exterior routing protocol that guarantees the loop-free exchange of routing information between autonomous systems. The Alcatel-Lucent implementation of BGP is designed for enterprise networks, specifically for border routers handling a public network connection, such as the organization's Internet Service Provider (ISP) link.

BGP Graceful Restart

BGP Graceful Restart is now supported and is enabled by default. On OmniSwitch devices in a redundant CMM configuration, during a CMM takeover/failover, interdomain routing is disrupted. Alcatel-Lucent Operating System BGP needs to retain forwarding information and also help a peering router performing a BGP restart to support continuous forwarding for inter-domain traffic flows by following the BGP graceful restart mechanism. This implementation supports BGP Graceful Restart mechanisms as defined in the RFC 4724.

DHCP / UDP Relay

DHCP Relay allows for forwarding of DHCP broadcast requests to configurable DHCP server IP address in a routing environment.

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.

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.

Per-VLAN DHCP Relay

It is possible to configure multiple DHCP relay (ip helper) addresses on a per-vlan basis. For the Per- VLAN service, identify the number of the VLAN that makes the relay request. You may identify one or more server IP addresses to which DHCP packets will be sent from the specified VLAN. Both standard and per VLAN modes are supported.

UDP Relay

In addition to BOOTP/DHCP relay, generic UDP relay is available. Using generic UDP relay, traffic destined for well-known service ports (e.g., NBNS/NBDD, DNS, TFTP) or destined for a user-defined service port can be forwarded to specific VLANs on the switch.

DNS Client

A Domain Name System (DNS) resolver is an internet service that translates host names into IP addresses. Every time you enter a host name, a DNS service must look up the name on a server and resolve the name to an IP address. You can configure up to three domain name servers that will be queried in turn to resolve the host name. If all servers are queried and none can resolve the host name to an IP address, the DNS fails. If the DNS fails, you must either enter an IP address in place of the host name or specify the necessary lookup tables on one of the specified servers.

Generic Routing Encapsulation

Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels. GRE is used to create a virtual point-to-point link between routers at remote points in a network. This feature supports the creation, administration, and deletion of IP interfaces whose underlying virtual device is a GRE tunnel.

IP Multinetting

IP multinetting allows multiple subnets to coexist within the same VLAN domain. This implementation of the multinetting feature allows for the configuration of up to eight IP interfaces per a single VLAN. Each interface is configured with a different subnet.

IP Route Map Redistribution

Route map redistribution provides the ability to control which routes from a source protocol are learned and distributed into the network of a destination protocol. A route map consists of one or more user-defined statements that can determine which routes are allowed or denied access to the network. In addition, a route map may also contain statements that modify route parameters before they are redistributed.

Redistribution is configured by specifying a source and destination protocol and the name of an existing route map. Criteria specified in the route map is applied to routes received from the source protocol.

IP-IP Tunneling

The IP/IP tunneling feature allows IP traffic to be tunneled through an IP network. This feature can be used to establish connectivity between remote IP networks using an intermediate IP network such as the Internet.

OSPFv2

OSPF is a shortest path first (SPF), or link-state, protocol for IP networks. Also considered an interior gateway protocol (IGP), it distributes routing information between routers in a single Autonomous System (AS). OSPF chooses the least-cost path as the best path. OSPF is suitable for complex networks with a large number of routers by providing faster convergence, loop free routing, and equal-cost multi-path routing where packets to a single destination can be sent to more than one interface simultaneously. OSPF adjacencies over non-broadcast links are also supported.

In addition, OSPFv2 supports graceful (hitless) support during failover, which is the time period between the restart and the reestablishment of adjacencies after a planned (e.g., the users performs the takeover) or unplanned (e.g., the primary management module unexpectedly fails) failover.

RIPv1/RIPv2

Routing Information Protocol (RIP) is a widely used Interior Gateway Protocol (IGP) that uses hop count as its routing metric. RIP-enabled routers update neighboring routers by transmitting a copy of their own routing table. The RIP routing table uses the most efficient route to a destination, that is, the route with the fewest hops and longest matching prefix.

The OmniSwitch supports RIP version 1 (RIPv1), RIP version 2 (RIPv2), and RIPv2 that is compatible with RIPv1. In addition, text key and MD5 authentication, on an interface basis, for RIPv2 is also supported as well as ECMP for up to 16 paths.

RIP Timer Configuration

Update—The time interval between advertisement intervals.

Invalid—The amount of time before an active route expires and transitions to the garbage state.

Garbage—The amount of time an expired route remains in the garbage state before it is removed from the RIB.

Holddown—The amount of time during which a route remains in the hold-down state.

Routing Protocol Preference

Specifying a routing protocol preference is supported. This is done by configuring a weight for each routing protocol (including static routes) to control which entry to prefer when two entries exist from different sources.

Server Load Balancing (SLB)

Server Load Balancing (SLB) software provides a method to logically manage a group of physical servers sharing the same content (known as a server farm) as one large virtual server (known as an SLB cluster). SLB clusters are identified and accessed at Layer 3 by the use of Virtual IP (VIP) addresses or at Layer 2 or Layer 3 by the use of a QoS policy condition. The OmniSwitch operates at wire speed to process client requests addressed to the VIP of an SLB cluster or classified by a QoS policy condition and send them to the physical servers within the cluster.

Using SLB clusters can provide cost savings (costly hardware upgrades can be delayed or avoided), scalability (as the demands on your server farm grow you can add additional physical servers), reliability (if one physical server goes down the remaining servers can handle the remaining workload), and flexibility (you can tailor workload requirements individually to servers within a cluster).

Server Load Balancing - WRR

Enhances the Server Load Balancing to allow for the configuration of a Weighted Round Robin distribution algorithm. When configured, SLB will distribute traffic according to the relative “weight” a server has within an SLB cluster.

VRRPv2

VRRP is a standard router redundancy protocol that provides redundancy by eliminating the single point of failure inherent in a default route environment. VRRP allows for the configuration of a virtual router called the master router, and is responsible for forwarding virtual router advertisements. If the master router becomes unavailable, the highest priority backup router will transition to the master state.

VRRP allows routers on a LAN to back up a static default route with a virtual router. VRRP dynamically assigns responsibility for a virtual router to a physical router (VRRP router) on the LAN. The virtual router is associated with an IP address (or set of IP addresses) on the LAN. A virtual router master is elected to forward packets for the virtual router’s IP address. If the master router becomes unavailable, the highest priority backup router will transition to the master state.

VRRP supports VRRP Tracking. A virtual router’s priority may be conditionally modified to prevent another router from taking over as master. Tracking policies are used to conditionally modify the priority setting whenever an IP interface, slot/port, and/or IP address associated with a virtual router goes down.

IPv6 Feature Support

IPv6 is designed as a successor to IPv4. The changes from IPv4 to IPv6 fall primarily into the following categories:

- Address size increased from 32 bits (IPv4) to 128 bits (IPv6)
- Dual Stack IPv4/IPv6
- ICMPv6
- Neighbor Discovery
- Stateless Autoconfiguration
- OSPFv3
- RIPng
- Static Routes
- Tunneling: Configured and 6-to-4 dynamic tunneling
- Ping6, Traceroute6
- DNS client using Authority records
- Telnetv6 - Client and server
- FTPv6 – Client and server
- SSHv6 – Client and Server

Globally Unique Local Unicast Addresses

Unique Local IPv6 Unicast Addresses are intended to be routable within a limited area such as a site but not on the global Internet. Unique Local IPv6 Unicast Addresses are used in conjunction with BGP (IBGP) speakers as well as exterior BGP (EBGP) neighbors based on configured policies and have the following characteristics:

- Globally unique ID (with high probability of uniqueness).
- Use the well-known prefix FC00::/7 to allow for easy filtering at site boundaries.
- Allow sites to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces that use these prefixes.
- Internet Service Provider independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity.
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses.
- In practice, applications may treat these addresses like global scoped addresses.
- A 40-bit global identifier is used to make the local IPv6 address prefixes globally unique. This global ID can either be explicitly configured, or created using the pseudo-algorithm recommended in RFC 4193.

Scoped Multicast Addresses

The IPv6 Scoped Multicast Address feature allows for the configuration of per-interface scoped IPv6 multicast boundaries. This feature allows an OmniSwitch to configure a PIM domain into multiple administratively scoped regions and is known as a Zone Boundary Router (ZBR). A ZBR will not forward packets matching an interface's boundary definition into or out of the scoped region, will prune the boundary for PIM-DM, as well as reject joins for the scoped range for PIM-SM.

BGP4

The Border Gateway Protocol (BGP) is an exterior routing protocol that guarantees the loop-free exchange of routing information between autonomous systems. The Alcatel-Lucent implementation of BGP is designed for enterprise networks, specifically for border routers handling a public network connection, such as the organization's Internet Service Provider (ISP) link.

BGP IPv6 Extensions

The Omniswitch provides IPv6 support for BGP using Multiprotocol Extensions. The same procedures used for IPv4 prefixes can be applied for IPv6 prefixes as well and the exchange of IPv4 prefixes will not be affected by this feature. However, there are some attributes that are specific to IPv4, such as AGGREGATOR, NEXT_HOP and NLRI. Multiprotocol Extensions for BGP also supports backward compatibility for the routers that do not support this feature. This implementation supports Multiprotocol BGP as defined in the following RFCs 4760 and 2545.

IPsec Support for IPv6, OSPFv3, RIPng

IPsec is a suite of protocols for securing IPv6 communications by authenticating and/or encrypting each IPv6 packet in a data stream. IPsec provides security services such as encrypting traffic, integrity validation, authentication, and anti-replay.

The OmniSwitch implementation of IPsec supports the transport mode of operation and manually configured SAs only. In transport mode, the data transferred (payload) in the IPv6 packet is encrypted and/or authenticated and only the payloads that are originated and destined between two end-points are processed with IPsec.

OSPFv3

OSPFv3 is an extension of OSPF version 2 (OSPFv2) that provides support for networks using the IPv6 protocol. OSPFv2 is for IPv4 networks.

Both versions of OSPF are shortest path first (SPF), or link-state, protocols for IP networks. Also considered interior gateway protocols (IGP), both versions distribute routing information between routers in a single Autonomous System (AS). OSPF chooses the least-cost path as the best path. OSPF is suitable for complex networks with a large number of routers by providing faster convergence, loop free routing, and equal-cost multi-path routing where packets to a single destination can be sent to more than one interface simultaneously. OSPF adjacencies over non-broadcast links are also supported.

RIPng

The OmniSwitch supports Routing Information Protocol next generation (RIPng) for IPv6 networks. RIPng is based on RIPv1/RIPv2 and is an Interior Gateway Protocol (IGP) best suited for moderate sized networks.

VRRPv2/VRRPv3

Similar to VRRPv2, VRRPv3 is a standard router redundancy protocol that provides redundancy by eliminating the single point of failure inherent in a default route environment. The VRRPv3 router, which controls the IPv6 address associated with a virtual router is called the master router, and is responsible for forwarding virtual router advertisements. If the master router becomes unavailable, the highest priority backup router will transition to the master state.

Both versions of VRRP allow routers on a LAN to back up a static default route with a virtual router. VRRP dynamically assigns responsibility for a virtual router to a physical router (VRRP router) on the LAN. The virtual router is associated with an IP address (or set of IP addresses) on the LAN. A virtual router master is elected to forward packets for the virtual router's IP address. If the master router becomes unavailable, the highest priority backup router will transition to the master state.

In addition, both versions support VRRP Tracking. A virtual router's priority may be conditionally modified to prevent another router from taking over as master. Tracking policies are used to conditionally modify the priority setting whenever an ip interface, slot/port, and/or IP address associated with a virtual router goes down.

QoS Feature Support

The OmniSwitch software and queue management architecture provide a way to identify traffic entering the network and manipulate flows coming through the switch. The flow manipulation (generally referred to as Quality of Service or QoS) can be as simple as configuring QoS policies to allow/deny traffic or as complicated as remapping 802.1p bits from a Layer 2 network to ToS values in a Layer 3 network.

The types of policies typically used include, but are not limited to, the following:

- Basic QoS—includes traffic prioritization and bandwidth shaping.
- ICMP policies—includes filtering, prioritizing, and/or rate limiting ICMP traffic for security.
- 802.1p/ToS/DSCP—includes policies for marking and mapping including support for entering a range of DSCP values.
- Policy Based Routing (PBR)—includes policies for redirecting routed traffic.
- Policy Based Mirroring—includes mirror-to-port (MTP) policies for mirroring ingress, egress, or both ingress and egress traffic.
- Access Control Lists (ACLs)—ACLs are a specific type of QoS policy that is used for Layer 2 and Layer 3/4 filtering.

This implementation of QoS integrates traffic management with QoS scheduling. Embedded profiles apply the QoS admission control and bandwidth management configurations to traffic flows. Packets received by the switch are classified on the ingress and queue management is applied on the egress to avoid congestion.

Auto-QoS Prioritization for NMS Traffic

This feature can be used to enable the automatic prioritization of NMS traffic—SSH (TCP Port 22), Telnet (TCP Port 23), WebView (HTTP Port 80) and SNMP (TCP port 161)—that is destined for the switch. Prioritization maximizes access for NMS traffic and helps to reduce the potential for DoS attacks.

Ingress and Egress Bandwidth Shaping

Bandwidth shaping is configured on a per port basis by specifying a maximum bandwidth value for ingress and egress ports.

Policy Based Routing (Permanent Mode)

Policy Based Routing may be used to redirect traffic to a particular gateway based on source or destination IP address, source or destination network group, source or destination TCP/UDP port, a service or service group, IP protocol, or built-in source port group.

Traffic may be redirected to a particular gateway regardless of what routes are listed in the routing table. Note that the gateway address does not have to be on a directly connected VLAN; the address may be on any network that is learned by the switch.

Policy Lists

A policy list is a group of policy rules that is identified by the list name. There are two types of lists available:

Default—All rules are associated with a default policy list when the rules are created. This list is not configurable, but it is possible to direct QoS not to assign a rule to this list. Default policy list rules are applied to ingress traffic.

Universal Network Profile (UNP)—This type of policy list is associated with a UNP. The rules in this list are applied to ingress traffic that is classified into the user profile. Up to 13 policy lists (including the default list) are supported per switch. Only one policy list per UNP is allowed, but a policy list can be associated with multiple profiles.

Redirect Policies (Port and Link Aggregate)

Two policy action commands are available for configuring QoS redirection policies: policy action redirect port and policy action redirect linkagg. A redirection policy sends traffic that matches the policy to a specific port or

link aggregate instead of the originally intended destination. This type of policy may use any condition; the policy action determines which port or link aggregate to which the traffic is sent.

Tri-Color Marking

Tri-Color Marking (TCM) provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. The TCM policer meters traffic based on user-configured packet rates and burst sizes and then marks the metered packets as green, yellow, or red based on the metering results.

TCM policer meters each packet and passes the metering result along with the packet to the Marker. Depending upon the result sent by the Meter, the packet is then marked with either the green, yellow, or red color. The marked packet stream is then transmitted on the egress based on the color-coded priority assigned.

The TCM Meter operates in Color-Blind mode (the Color-Aware mode is not supported). In the Color- Blind mode, the Meter assumes that the incoming packet stream is uncolored. However incoming packets with the CFI/DEI bit set are automatically given an internal lower priority.

There are two types of TCM marking supported:

- **Single-Rate TCM (srTCM) according to RFC 2697**—Packets are marked based on a Committed Information Rate (CIR) and two associated burst size values: Committed Burst Size (CBS) and Peak Burst Size (PBS).
- **Two-Rate TCM (trTCM) according to RFC 2698**—Packets are marked based on a CIR value *and* a Peak Information Rate (PIR) value and two associated burst size values: CBS and PBS.

Both srTCM and trTCM handle the burst in the same manner. The main difference between the two types is that srTCM uses one rate limiting value (CIR) and trTCM uses two rate limiting values (CIR and PIR) to determine packet marking.

Multicast Feature Support

DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) is a dense-mode multicast routing protocol. DVMRP—which is essentially a “broadcast and prune” routing protocol—is designed to assist routers in propagating IP multicast traffic through a network. DVMRP works by building per-source broadcast trees based on routing exchanges, then dynamically creating per-source, group multicast delivery trees by pruning the source’s truncated broadcast tree.

IGMP Multicast Group Configuration Limit

By default there is no limit on the number of IGMP groups that can be learned on a port/VLAN instance. However, a user can now configure a maximum group limit to limit the number of IGMP groups that can be learned. The maximum group limit can be applied globally, per VLAN, or per port. Port settings override VLAN settings, which override global settings. Once the limit is reached, the user can configure the switch to drop the incoming membership request, or replace an existing membership with the incoming membership request. This feature is available on IPv4 and IPv6/MLD.

IGMP Relay - Relay IGMP Packets to Specific Host

Encapsulates unicast IGMP packets to the specified multicast server. This immediately notifies the multicast server to forward a new multicast stream when a subscriber has joined the new group without relying on the L3 multicast network (e.g. PIM) to propagate this event.

IP Multicast Switching (IPMS) – IPv4/IPv6

IP Multicast Switching is a one-to-many communication technique employed by emerging applications such as video distribution, news feeds, conferencing, netcasting, and resource discovery (OSPF, RIP2, and BOOTP). Unlike unicast, which sends one packet per destination, multicast sends one packet to all devices in any subnetwork that has at least one device requesting the multicast traffic. Multicast switching also requires much less bandwidth than unicast techniques and broadcast techniques since the source hosts only send one data stream to the ports on which destination hosts that request it are attached.

Destination hosts signal their intent to receive a specific multicast stream by sending a request to do so to a nearby switch using Internet Group Management Protocol (IGMP). The switch then learns on which ports multicast group subscribers are attached and can intelligently deliver traffic only to the respective ports. This mechanism is often referred to as IGMP snooping (or IGMP glean). Alcatel-Lucent’s implementation of IGMP snooping is called IP Multicast Switching (IPMS). IPMS allows switches to efficiently deliver multicast traffic in hardware at wire speed.

Both IGMP version 3 (IGMPv3), which handles forwarding by source IP address and IP multicast destination, and IGMP version 2 (IGMPv2), which handles forwarding by IP multicast destination address only, are supported.

IP Multicast Switching (IPMS) - Proxying

IP multicast proxying and configuring the IGMP and MLD unsolicited report interval are available with this implementation of IPMS. Proxying enables the aggregation of IGMP and MLD group membership information and the reduction in reporting queriers. The unsolicited report interval refers to the time period in which to proxy any changed IGMP membership state.

L2 Static Multicast Addresses

Static multicast MAC addresses are used to send traffic intended for a single destination multicast MAC address to multiple switch ports within a given VLAN. A static multicast address is assigned to one or more switch ports for a given VLAN. The ports associated with the multicast address are then identified as egress ports. When traffic received on ports within the same VLAN is destined for the multicast address, the traffic is forwarded on the egress ports that are associated with the multicast address.

One of the benefits of using static multicast addresses is that multicast traffic is switched in hardware and no longer subject to flood limits on broadcast traffic.

PIM-SM/PIM-DM/PIM-SSM

Protocol-Independent Multicast (PIM) is an IP multicast routing protocol that uses routing information provided by unicast routing protocols, such as RIP and OSPF. PIM is “protocol-independent” because it does not rely on any particular unicast routing protocol. Sparse mode PIM (PIM-SM) contrasts with flood-and-prune dense mode multicast protocols, such as DVMRP and PIM Dense Mode (PIM-DM) in that multicast forwarding in PIM-SM is initiated only via specific requests, referred to as Join messages.

PIM-DM for IPv4 is supported. PIM-DM packets are transmitted on the same socket as PIM-SM packets, as both use the same protocol and message format. Unlike PIM-SM, in PIM-DM there are no periodic joins transmitted; only explicitly triggered prunes and grafts. In addition, there is no Rendezvous Point (RP) in PIM-DM.

Protocol Independent Multicast Source-Specific Multicast (PIM-SSM) is a highly-efficient extension of PIM. SSM, using an explicit channel subscription model, allows receivers to receive multicast traffic directly from the source; an RP tree model is not used. In other words, a Shortest Path Tree (SPT) between the receiver and the source is created without the use of a Rendezvous Point (RP).

Monitoring and Troubleshooting Feature Support

DDM - Digital Diagnostic Monitoring

Digital Diagnostics Monitoring allows an OmniSwitch to monitor the status of a transceiver by reading the information contained on the transceiver's EEPROM. The transceiver can display Actual, Warning-Low, Warning-High, Alarm-Low and Alarm-High for the following:

- Temperature
- Supply Voltage
- Current
- Output (Transmit) Power
- Input (Receive) Power
- Traps can be enabled if any of these above values crosses the pre-defined low or high thresholds of the transceiver.
- **Note:** Not all transceivers support DDM, refer to the Transceivers Guide for additional DDM information.

Health Statistics

To monitor resource availability, the NMS (Network Management System) needs to collect significant amounts of data from each switch. As the number of ports per switch (and the number of switches) increases, the volume of data can become overwhelming. The Health Monitoring feature can identify and monitor a switch's resource utilization levels and thresholds, improving the efficiency in data collection. Health Monitoring provides the following data to the NMS:

- Switch-level input/output, memory and CPU utilization levels
- Module-level and port-level input/output utilization levels

For each monitored resource, the following variables are defined:

- Most recent utilization level (percentage)
- Average utilization level over the last minute (percentage)
- Average utilization level over the last hour (percentage)
- Maximum utilization level over the last hour (percentage)
- Threshold level

Additionally, Health Monitoring provides the capacity to specify thresholds for the resource utilization levels it monitors, and generates traps based on the specified threshold criteria.

Ping and Traceroute

Ping and Traceroute support both IPv4 and IPv6 along with additional parameters such as a source interface and timeout.

Port Mirroring

Port mirroring allows transmitted and received traffic from a “mirrored” port to be copied to another port. The “mirroring” port receives a copy of all transmitted and received traffic and can be used to send the traffic to a network analyzer.

Port Mirroring – Policy-Based

This feature enhances the port mirroring functionality on the OmniSwitch. It allows policies to be configured to determine when traffic should be mirrored based on policies rather than being restricted to a specified port. The following policies can be configured:

- Traffic between 2 ports
- Traffic from a source address
- Traffic to a destination address
- Traffic to/from an address
- Traffic between 2 addresses
- Traffic with a classification criterion based on packet contents other than addresses (for example , based on protocol, priority).
- VLAN-based mirroring - mirroring of packets entering a VLAN.
- Policy-Based Mirroring limitations:
 - The policy mirror action must specify the same analyzer port for all policies in which the action is used.
 - One policy-based mirroring session supported per switch.
 - One port-based mirroring session supported per switch. Note that policy-based and port-base mirroring are both allowed on the same port at the same time.
 - One remote port-based mirroring session supported per switch.
 - One port-monitoring session supported per switch.

Port Mirroring – Remote (802.1Q Based)

This feature provides a remote port mirroring capability where traffic from a local port can be carried across the network to an egress port where a sniffer can be attached. This feature makes use of an 802.1q tag to send the mirrored traffic over the network using tagged VLANs.

- There must not be any physical loop present in the remote port mirroring VLAN.
- Spanning Tree must be disabled for the remote port mirroring VLAN.
- Source learning must be disabled or overridden on the ports belonging to the remote port mirroring VLAN on the intermediate and destination switches.

Port Monitoring

The Port Monitoring feature allows you to examine packets to and from a specific Ethernet port (either ingress or egress) and capture the output to a file. Once a file is captured, you can FTP it to a Protocol Analyzer or PC for viewing.

RMON

Remote Network Monitoring (RMON) is an SNMP protocol used to manage networks remotely. RMON probes can be used to collect, interpret, and forward statistical data about network traffic from designated active ports in a LAN

segment to an NMS (Network Management System) application for monitoring and analyzing without negatively impacting network performance. RMON software is fully integrated in the software to acquire statistical information.

sFlow

sFlow is a network monitoring technology that gives visibility to the activity of the network, by providing network usage information. It provides the data required to effectively control and manage the network usage. sFlow is a sampling technology that meets the requirements for a network traffic monitoring solution.

sFlow is a sampling technology embedded within switches/routers. It provides the ability to monitor the traffic flows. It requires an sFlow agent software process running as part of the switch software and an sFlow collector, which receives and analyses the monitored data. The sFlow collector makes use of SNMP to communicate with an sFlow agent in order to configure sFlow monitoring on the device (switch).

Switch Logging

The Switch Logging feature is designed to provide a high-level event logging mechanism that can be useful in maintaining and servicing the switch. Switch Logging uses a formatted string mechanism to process log requests from applications. When a log request is received, Switch Logging verifies whether the Severity Level included with the request is less than or equal to the Severity Level stored for the appropriate Application ID. If it is, a log message is generated using the formatting specified by the log request and placed on the Switch Log Queue, and Switch Logging returns control back to the calling application. Otherwise, the request is discarded. The default output device is the log file located in the Flash File System. Other output devices can be configured via Command Line Interface. All log records generated are copied to all configured output devices.

Command Line Interface can be used to display and configure Switch Logging information. Log information can be helpful in resolving configuration or authentication issues, as well as general errors.

Metro Ethernet Feature Support

Ethernet Ring Protection (ERP) – G.8032

Ethernet Ring Protection (ERP) switching is a self-configuring algorithm that maintains a loop-free topology while providing data path redundancy and network scalability. ERP provides fast recovery times for Ethernet ring topologies by utilizing traditional Ethernet MAC and bridge functions.

This implementation of ERP is based on ITU-T G.8032 and uses the ring Automatic Protection Switching (APS) protocol to coordinate the prevention of network loops within a bridged Ethernet ring. Loop prevention is achieved by allowing the traffic to flow on all but one of the links within the protected Ethernet ring. This link is blocked and is referred to as the Ring Protection Link (RPL). When a ring failure condition occurs, the RPL is unblocked to allow the flow of traffic to continue through the ring.

- **Overlapping Protected VLANs on a Single Node**

In a network where all connected nodes cannot belong to a single ERP ring, the OmniSwitch supports multiple ERP rings. Each of the ERP rings has a different Service VLAN configured which allows the ERP PDUs to be processed by the corresponding ERP ring nodes. The Service VLANs configured for each of the ERP rings can be configured as a protected VLAN on the other ERP ring. The protected VLANs can be shared across ERP rings.

Ethernet Services

Ethernet Services provides a mechanism for tunneling multiple customer VLANs (CVLAN) through a service provider network over the Ethernet Metropolitan Area Network (EMAN). The service provider network uses one or more service provider VLANs (SVLAN) by appending an 802.1Q double tag or VLAN Translation on a customer port that contains the customer's assigned tunnel ID. This traffic is then encapsulated into the tunnel and transmitted through the service provider network. It is received on another Provider Edge (PE) that has the same tunnel ID.

This feature enables service providers to provide their customers with Transparent LAN Services (TLS). This service is multipoint in nature so as to support multiple customer sites or networks distributed over the edges of a service provider network.. Ethernet Services provides the following:

Ethernet service-based approach that is similar to configuring a virtual private LAN service (VPLS).

Ingress bandwidth sharing across User Network Interface (UNI) ports.

Ingress bandwidth rate limiting on a per UNI port, per CVLAN, or CVLAN per UNI port basis.

CVLAN (inner) tag 802.1p-bit mapping to SVLAN (outer) tag 802.1p bit.

CVLAN (inner) tag DSCP mapping to SVLAN (outer) tag 802.1p bit.

Profiles for saving and applying traffic engineering parameter values.

Capability to suspend the use of SAP bandwidth and priority actions allowing QoS rules for advanced classification of SAP traffic, such as mapping several DSCP/ToS values to the same outer 802.1p value.

Ethernet Services - Egress Rate Limiting

This feature allows for egress rate limiting for traffic going out on UNI ports. When a SAP is configured and bound to a SAP profile, the following information is used to provide egress rate limiting on traffic going out on the UNI port

Destination port = UNI port defined in the sap

VLAN = CVLAN defined in the sap (could be untagged, cvlan all or specific vlan id)

Rate limiter with the sap-profile egress-bandwidth

This feature does not support egress-rate limiting on IPMVLAN.

Ethernet Services - Tunneling L2 Protocols

Enhances the User Network Interface (UNI) profile to allow the control packets for 802.1x, 802.1ab, 802.3ad, 802.3ah, MVRP, STP and AMAP to be tunneled, discarded, or peered on UNI ports.

Note: 802.3ad and 802.3ah packets use the same MAC address. Therefore, the configuration for 802.3ad also applies to 802.3ah control packets.

Security Feature Support

Access Control Lists (ACLs)

Access Control Lists (ACLs) are Quality of Service (QoS) policies used to control whether or not packets are allowed or denied at the switch or router interface. ACLs are sometimes referred to as filtering lists. ACLs are distinguished by the kind of traffic they filter. In a QoS policy rule, the type of traffic is specified in the policy condition. The policy action determines whether the traffic is allowed or denied. In general, the types of ACLs include:

Layer 2 ACLs—for filtering traffic at the MAC layer. Usually uses MAC addresses or MAC groups for filtering.

Layer 3/4 ACLs—for filtering traffic at the network layer. Typically uses IP addresses or IP ports for filtering; note that IPX filtering is not supported.

Multicast ACLs—for filtering IGMP traffic.

ICMP drop rules—Allows condition combinations in policies that will prevent user pings, thus reducing DoS exposure from pings. Two condition parameters are also available to provide more granular filtering of ICMP packets: `icmptype` and `icmpcode`.

TCP connection rules—Allows the determination of an established TCP connection by examining TCP flags found in the TCP header of the packet. Two condition parameters are available for defining a TCP connection ACL: `established` and `tcpflags`.

Early ARP discard—ARP packets destined for other hosts are discarded to reduce processing overhead and exposure to ARP DoS attacks. No configuration is required to use this feature, it is always available and active on the switch. Note that ARPs intended for use by a local subnet, AVLAN, and VRRP are not discarded.

UserPorts—A port group that identifies its members as user ports to prevent spoofed IP traffic. When a port is configured as a member of this group, packets received on the port are dropped if they contain a source IP network address that does not match the IP subnet for the port.

UserPorts Profile—In addition to spoofed traffic, it is also possible to configure a global UserPorts profile to specify additional types of traffic, such as BPDU, RIP, OSPF, DVMRP, PIM, DHCP server response packets, DNS and/or BGP, to monitor on user ports. The UserPorts profile also determines whether user ports will filter the unwanted traffic or will administratively shutdown when the traffic is received. Note that this profile only applies to those ports that are designated as members of the UserPorts port group.

DropServices—A service group that improves the performance of ACLs that are intended to deny packets destined for specific TCP/UDP ports. This group only applies to ports that are members of the UserPorts group. Using the DropServices group for this function minimizes processing overhead, which otherwise could lead to a DoS condition for other applications trying to use the switch.

Access Control Lists (ACLs) for IPv6

Support for IPv6 ACLs on the OmniSwitch available. The following QoS policy conditions are available for configuring ACLs to filter IPv6 traffic. Note the following when using IPv6 ACLs:

Trusted/untrusted behavior is the same for IPv6 traffic as it is for IPv4 traffic.

IPv6 policies do not support the use of network groups, service groups, map groups, or MAC groups.

IPv6 multicast policies are not supported.

Anti-spoofing and other UserPorts profiles/filters do not support IPv6.

The default (built-in) network group, “Switch”, only applies to IPv4 interfaces. There is no such group for IPv6 interfaces.

Account and Password Policies

This feature allows a switch administrator to configure password policies for password creation and management. The administrator can configure how often a password must be changed, lockout settings for failed attempts, password complexity, history, and age as well as other account management settings.

Admin User Remote Access Control

The OmniSwitch can be configured to allow the admin user to only have access to the switch via the console port.

ARP Defense Optimization

This feature enhances how the OmniSwitch can respond to an ARP DoS attack by not adding entries to the forwarding table until the net hop ARP entry can be resolved.

ARP Poisoning Detect

This feature detects the presence of an ARP-Poisoning host on the network using configured restricted IP addresses for which the switch, on sending an ARP request, should not get back an ARP response. If an ARP response is received, the event is logged and the user is alerted using an SNMP trap.

By default ARP requests are not added to the ARP cache. Only router solicited ARP requests will be added to the cache.

Authenticated Switch Access

Authenticated Switch Access (ASA) is a way of authenticating users who want to manage the switch. With authenticated access, all switch login attempts using the console or modem port, Telnet, FTP, SNMP, or HTTP require authentication via the local user database or via a third-party server. The type of server may be an authentication-only mechanism or an authentication, authorization, and accounting (AAA) mechanism.

AAA servers are able to provide authorization for switch management users as well as authentication. (They also may be used for accounting.) User login information and user privileges may be stored on the servers. The following AAA servers are supported on the switch:

- Remote Authentication Dial-In User Service (RADIUS). Authentication using this type of server was certified with Juniper Steel Belted RADIUS server (any industry standard RADIUS server should work).

- Lightweight Directory Access Protocol (LDAP).

By default, switch management users may be authenticated through the console port via the local user database. If external servers are configured for other management interfaces but the servers become unavailable, the switch will poll the local user database for login information if the switch is configured for local checking of the user database. The database includes information about whether or not a user is able to log into the switch and what kinds of privileges or rights the user has for managing the switch.

IP DoS Filtering

By default, the switch filters the following denial of service (DoS) attacks, which are security attacks aimed at devices that are available on a private network or the Internet:

- ARP Flood Attack

- Invalid IP Attack

- Multicast IP and MAC Address Mismatch

- Ping Overload

- Packets with loopback source IP address

Learned Port Security (LPS)

Learned Port Security (LPS) provides a mechanism for authorizing source learning of MAC addresses on 10/100/1000, Gigabit, and Gigabit Ethernet ports. Using LPS to control source MAC address learning provides the following benefits:

- A configurable source learning time limit that applies to all LPS ports.
 - A configurable limit on the number of MAC addresses allowed on an LPS port.
 - Dynamic configuration of a list of authorized source MAC addresses.
 - Static configuration of a list of authorized source MAC addresses.
 - Two methods for handling unauthorized traffic: Shutting down the port or only blocking traffic that violates LPS criteria.
 - A configurable limit to the number of filtered MAC addresses allowed on an LPS port. Conversion of dynamically learned MAC addresses to static MAC address entries.
 - Support for all authentication methods and LPS on the same switch port.
- LPS has the following limitations:
- You cannot configure LPS on link aggregate ports.

Learned MAC Address Notification

The LPS feature enables the OmniSwitch to generate an SNMP trap when a new bridged MAC address is learned on an LPS port. A configurable trap threshold number is provided to determine how many MAC addresses are learned before such traps are generated for each MAC address learned thereafter. Trap contents includes identifying information about the MAC, such as the address itself, the corresponding IP address, switch identification, and the slot and port number on which the MAC was learned.

Policy Server Management

Policy servers use Lightweight Directory Access Protocol (LDAP) to store policies that are configured through Alcatel-Lucent's PolicyView network management application. PolicyView is an OmniVista application that runs on an attached workstation.

The Lightweight Directory Access Protocol (LDAP) is a standard directory server protocol. The LDAP policy server client in the switch is based on RFC 2251. Currently, PolicyView is supported for policy management.

Port Mapping (Private VLANs)

Port Mapping is a security feature that controls peer users from communicating with each other. A Port Mapping session comprises a session ID and a set of user ports and/or a set of network ports. User ports within a session cannot communicate with each other and can only communicate via network ports. In a Port Mapping session with user port set A and network port set B, ports in set A can only communicate with ports in set B. If set B is empty, ports in set A can communicate with rest of the ports in the system.

A port mapping session can be configured in unidirectional or bidirectional mode. In the unidirectional mode, the network ports can communicate with each other within the same session. In the bidirectional mode, the network ports cannot communicate with each other. Network ports of a unidirectional port mapping session can be shared with other unidirectional sessions, but cannot be shared with any sessions configured in bidirectional mode. Network Ports of different sessions can communicate with each other.

SNMP Traps

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

No.	Trap Name	Platforms	Description
0	coldStart	OS10K OS6900	The SNMP agent in the switch is reinitiating and itsk configuration may have been altered.
1	warmStart	OS10K OS6900	The SNMP agent in the switch is reinitiating itself and its configuration is unaltered.
2	linkDown	OS10K OS6900	The SNMP agent in the switch recognizes a failure in one of the communications links configured for the switch.
3	linkUp	OS10K OS6900	The SNMP agent in the switch recognizes that one of the communications links configured for the switch has come up.
4	authenticationFailure	OS10K OS6900	The SNMP agent in the switch has received a protocol message that is not properly authenticated.
5	entConfigChange	OS10K OS6900	An entConfigChange notification is generated when a conceptual row is created, modified, or deleted in one of the entity tables.
6	policyEventNotification	OS10K OS6900	The switch notifies the NMS when a significant event happens that involves the policy manager.
7	chassisTrapsStr	OS10K OS6900	A software trouble report (STR) was sent by an application encountering a problem during its execution.
8	chassisTrapsAlert	OS10K OS6900	A notification that some change has occurred in the chassis.
9	chassisTrapsStateChange	OS10K OS6900	An NI status change was detected.
10	chassisTrapsMacOverlap	OS10K OS6900	A MAC range overlap was found in the backplane eeprom.
11	vrrpTrapNewMaster	OS10K OS6900	The SNMP agent has transferred from the backup state to the master state.
12	vrrpTrapAuthFailure	OS10K OS6900	This trap is not supported.
13	healthMonModuleTrap	OS10K OS6900	Indicates a module-level threshold was crossed.
14	healthMonPortTrap	OS10K OS6900	Indicates a port-level threshold was crossed.
15	healthMonCmmTrap	OS10K OS6900	This trap is sent when the Module-level rising/falling threshold is crossed.
16	bgpEstablished	OS10K OS6900	The BGP routing protocol has entered the established state.
17	bgpBackwardTransition	OS10K OS6900	This trap is generated when the BGP router port has moved from a more active to a less active state.
18	esmDrvTrapDropsLink	OS10K OS6900	This trap is sent when the Ethernet code drops the link because of excessive errors.
19	portViolationTrap	OS10K	This trap is sent when a port violation occurs.

No.	Trap Name	Platforms	Description
		OS6900	The port violation trap will indicate the source of the violation and the reason for the violation.
20	dvmrpNeighborLoss	OS10K OS6900	A 2-way adjacency relationship with a neighbor has been lost. This trap is generated when the neighbor state changes from “active” to “one-way,” “ignoring” or “down.” The trap is sent only when the switch has no other neighbors on the same interface with a lower IP address than itself.
21	dvmrpNeighborNotPruning	OS10K OS6900	A non-pruning neighbor has been detected in an implementation-dependent manner. This trap is generated at most once per generation ID of the neighbor. For example, it should be generated at the time a neighbor is first heard from if the prune bit is not set. It should also be generated if the local system has the ability to tell that a neighbor which sets the prune bit is not pruning any branches over an extended period of time. The trap should be generated if the router has no other neighbors on the same interface with a lower IP address than itself.
22	risingAlarm	OS10K OS6900	An Ethernet statistical variable has exceeded its rising threshold. The variable’s rising threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
23	fallingAlarm	OS10K OS6900	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	OS10K OS6900	Sent by a bridge that became the new root of the spanning tree.
25	stpRootPortChange	OS10K OS6900	A root port has changed for a spanning tree bridge. The root port is the port that offers the lowest cost path from this bridge to the root bridge.
26	mirrorConfigError	OS10K OS6900	This trap is sent when any NI fails to configure mirroring. Due to this error, port mirroring session will be terminated.
27	mirrorUnlikeNi	OS10K OS6900	The mirroring configuration is deleted due to the swapping of different NI board type. The Port Mirroring session which was active on a slot cannot continue with the insertion of different NI type in the same slot.
28	slbTrapOperStatus	OS10K OS6900	A change occurred in the operational status of the server load balancing entity.
29	sessionAuthenticationTrap	OS10K OS6900	An authentication failure trap is sent each time a user authentication is refused.
30	trapAbsorptionTrap	OS10K OS6900	The absorption trap is sent when a trap has been absorbed at least once.
31	alaDoSTrap	OS10K	Indicates that the sending agent has received a

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

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

No.	Trap Name	Platforms	Description
		OS6900	bias current has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx bias current.
69	ddmTxPowerThresholdViolated	OS10K OS6900	This trap is sent when an SFP/ XFP/SFP+ Tx output power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx output power.
70	ddmRxPowerThresholdViolated	OS10K OS6900	This trap is sent when an SFP/ XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
71	webMgtServerErrorTrap	OS10K OS6900	This trap is sent to management station(s) when the Web Management server goes into error state after becoming unreachable twice within a minute.
72	multiChassisIpcVlanUp	OS10K OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is up.
73	multiChassisIpcVlanDown	OS10K OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is down.
74	multiChassisMisconfigurationFailure	OS10K OS6900	This trap is sent to indicate a mis-configuration due to Chassis Id or IPC VLAN.
75	multiChassisHelloIntervalConsisFailure	OS10K OS6900	This trap is sent to indicate a hello interval consistency failure.
76	multiChassisStpModeConsisFailure	OS10K OS6900	This trap is sent to indicate an STP mode consistency failure.
77	multiChassisStpPathCostModeConsisFailure	OS10K OS6900	This trap is sent to indicate an STP path cost mode consistency failure.
78	multiChassisVflinkStatusConsisFailure	OS10K OS6900	This trap is sent to indicate a VFLink status consistency failure.
79	multiChassisStpBlockingStatus	OS10K OS6900	This trap is sent to indicate the STP status for some VLANs on the VFLink is in blocking state.
80	multiChassisLoopDetected	OS10K OS6900	This trap is sent to indicate a loop has been detected.
81	multiChassisHelloTimeout	OS10K OS6900	This trap is sent to indicate the hello timeout has occurred.
82	multiChassisVflinkDown	OS10K OS6900	This trap is sent to indicate the VFLink is down.
83	multiChassisVFLMemberJoinFailure	OS10K	This trap is sent to indicate a port configured as

No.	Trap Name	Platforms	Description
		OS6900	virtual-fabric member is unable to join the virtual-fabric link.
84	alaDHLVlanMoveTrap	OS10K OS6900	When linkA or linkB goes down or comes up and both ports are part of some vlan-map, this trap is sent to the Management Entity, with the DHL port information.
85	alaDhcpClientAddressAddTrap	OS10K OS6900	This trap is sent when a new IP address is assigned to DHCP Cli-ent interface.
86	alaDhcpClientAddressExpiryTrap	OS10K OS6900	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address.
87	alaDhcpClientAddressModifyTrap	OS10K OS6900	This trap is sent when the DHCP client is unable to obtain the existing IP address and a new IP address is assigned to the DHCP client
88	vRtrIsisDatabaseOverload	Not supported	This notification is generated when the system enters or leaves the overload state.
89	vRtrIsisManualAddressDrops	Not supported	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
90	vRtrIsisCorruptedLSPDetected	Not supported	This notification is generated when an LSP that was stored in memory has become corrupted.
91	vRtrIsisMaxSeqExceedAttempt	Not supported	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
92	vRtrIsisIDLenMismatch	Not supported	A notification sent when a PDU is received with a different value of the System ID Length.
93	vRtrIsisMaxAreaAdrrsMismatch	Not supported	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
94	vRtrIsisOwnLSPPurge	Not supported	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
95	vRtrIsisSequenceNumberSkip	Not supported	When an LSP is received without a System ID and different contents.
96	vRtrIsisAutTypeFail	Not supported	A notification sent when a PDU is received with the wrong authentication type field.
97	vRtrIsisAuthFail	Not supported	A notification sent when a PDU is received with an incorrent authentication information field.
98	vRtrIsisVersionSkew	Not supported	A notification sent when a Hello PDU is received from an IS running a different version of the protocol.
99	vRtrIsisAreaMismatch	Not supported	A notification sent when a Hello PDU is received from an IS which does not share any area address.
100	vRtrIsisRejectedAdjacency	Not supported	A notification sent when a Hello PDU is received from an IS, but does not establish an adjacency due to a lack of resources.
101	vRtrIsisLSPTooLargeToPropagate	Not	A notification sent when an attempt to propagate

No.	Trap Name	Platforms	Description
		supported	an LSP which is larger than the dataLinkBlockSize for a circuit.
102	vRtrIsisOrigLSPBufSizeMismatch	Not supported	A notification sent when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for the originating L1LSP BufferSize or originating L2LSPBufferSize respectively. Also when a Level 1 LSP or Level2 LSP is received containing the originating LSPBufferSize option and the value in the PDU option field does not match the local value for originating L1LSP BufferSize or originatingL2LSP BufferSize respectively.
103	vRtrIsisProtoSuppMismatch	Not supported	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
104	vRtrIsisAdjacencyChange	Not supported	A notification sent when an adjacency changes state, entering or leaving state up. The first 6 bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
105	vRtrIsisCircIdExhausted	Not supported	A notification sent when ISIS cannot be started on a LAN interface because a unique circId could not be assigned due to the exhaustion of the circId space.
106	vRtrIsisAdjRestartStatusChange	Not supported	A notification sent when an adjacency's graceful restart status changes.
107	mvrpVlanLimitReachedEvent	OS10K OS6900	This trap is sent when the number of VLANs learned dynamically by MVRP reaches the configured limit.
108	alaHAVlanClusterPeerMismatch	OS10K OS6900	The trap is sent when parameter as configured for this cluster ID (Level 1 check) does not match across the MCLAG peers.
109	alaHAVlanMCPeerMismatch	OS10K OS6900	The trap is sent when when the cluster parameters are matching on the peers, but MCLAG is not configured or clusters are not in operational state.
110	alaHAVlanDynamicMAC	OS10K OS6900	The trap is sent when the dynamic MAC is learned on non-server cluster port
111	unpMcLagMacIgnored	OS10K OS6900	This trap is sent when a MAC/User is dropped because the VLAN does not exist or UNP is not enabled on the MCLAG.
112	unpMcLagConfigInconsistency	OS10K OS6900	This trap is sent when a configuration becomes "Out of Sync".
113	multiChassisGroupConsisFailure	OS10K OS6900	This trap is sent when there is an inconsistency between local and peer chassis group.
114	multiChassisTypeConsisFailure	OS10K OS6900	This trap is sent when there is an inconsistency between local and peer chassis group.
115	alaPimNonBidirHello	OS10K OS6900	This trap is sent when a bidir-capable router has received a PIM hello from a non-bidir-capable router. It is generated whenever the counter

No.	Trap Name	Platforms	Description
			alaPimsmNon-BidirHelloMsgsRcvd is incremented, subject to the rate limit specified by alaPimsmNonBidirHelloNotificationPeriod.
116	dot1agCfmFaultAlarm	OS10K OS6900	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	alaSaaPIterationCompleteTrap	OS10K OS6900	This trap is sent when an IP SAA iteration is completed.
118	alaSaaEthIterationCompleteTrap	OS10K OS6900	This trap is sent when when an eth-LB or Eth-DMM SAA iteration is completed.
119	alaSaaMacIterationCompleteTrap	OS10K OS6900	This trap is sent when a MAC iteration is complete.
120	virtualChassisStatusChange	OS10K OS6900	This trap is sent when a chassis status change is detected.
121	virtualChassisRoleChange	OS10K OS6900	This trap is sent when a chassis role change is detected.
122	virtualChassisVflStatusChange	OS10K OS6900	This trap is sent when s vflink status change is detected.
123	virtualChassisVflMemberPortStatusCh	OS10K OS6900	This trap is sent when a vflink member port has a change of status.
124	virtualChassisVflMemberPortJoinFail	OS10K OS6900	This trap is sent when a port configured as virtual-fabric member is unable to join the virtual-fabric link.
125	lldpRemTablesChange	OS10K OS6900	This trap is sent when the value of lldpStatsRemTablelastChange Time changes. It can be utilized by an NMS to trigger LLDP remote systems table maintenance polls.
126	vRtrLdpInstanceStateChange	OS10K OS6900	This trap is sent when the LDP module changes state either administratively or operationally.
127	evbFailedCdcPtlvTrap	OS10K OS6900	This trap is sent when bridge receives a CDCP packet with: <ul style="list-style-type: none"> - Wrong TLV type, or - Wrong OUI, or - Role is set to Bridge, or - Wrong default channel(scid), or - Incorrect channel number(scid).
128	evbFailedEvbTlvTrap	OS10K OS6900	This trap is sent when bridge receives an EVBTLV packet with: <ul style="list-style-type: none"> - Wrong TLV type. or - Incorrect TLV length, or - Wrong OUI.
129	evbUnknownVsiManagerTrap	OS10K OS6900	This trap is sent when bridge receives a VDP packet with: <ul style="list-style-type: none"> - Unknown Manager ID type, or - Wrong Manager ID length.
130	evbVdpAssocTlvTrap	OS10K	This trap is sent when bridge receives an

No.	Trap Name	Platforms	Description
		OS6900	ASSOC TLV in a VDP packet with: - Null VID found and number of entry field is not 1, or - Unknown filter format, - Null VID on De-Assoc TLV type, or - VSI included more than Max number of filter info entries
131	evbCdcplldpExpiredTrap	OS10K OS6900	This trap is sent when an LLDP Timer expired in bridge. The timer expires when LLDP doesn't not receive CDCP TLV within a specified interval.
132	evbTlvExpiredTrap	OS10K OS6900	This trap is sent when an LLDP Timer expired in bridge. The timer expires when LLDP doesn't not receive EVB TLV within a specified interval.
133	evbVdpKeepaliveExpiredTrap	OS10K OS6900	This trap is sent when a VDP Keep Alive Timer expired in bridge. The timer expires when the bridge doesn't not receive VDP Keep Alive message within a specified interval.
134	smgrServiceError	OS10K OS6900	This trap is sent when there is a failure to create/delete a service.
135	smgrServiceHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a service, or to program the hardware tables for a service.
136	smgrServiceSapError	OS10K OS6900	This trap is sent when there is a failure to create/delete a Service Access Point.
137	smgrServiceSapHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for a SAP, or to program the hardware tables for a SAP.
138	smgrServiceSdpError	OS10K OS6900	This trap is sent when there is a failure to create/delete a Service Distribution Point.
139	smgrServiceSdpHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP, or to program the hardware tables for an SDP.
140	smgrServiceSdpBindError	OS10K OS6900	This trap is sent when there is a failure to create/delete an SDP Bind.
141	smgrServiceSdpBindHwError	OS10K OS6900	This trap is sent when there is a failure to allocate/de-allocate a hardware resource for an SDP Bind, or to program the hardware tables for an SDP Bind.
142	smgrGeneralError	OS10K OS6900	This trap is sent when there is a .general system failure detected during normal system operation.
143	smgrStatusChange	OS10K OS6900	This trap is sent when there is a status change for a group of selected services.
144	portViolationNotificationTrap	OS10K	This trap is sent when a port violation is cleared.

No.	Trap Name	Platforms	Description
		OS6900	

Unsupported Software Features

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

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

Unsupported CLI Commands

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

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

Open Problem Reports and Feature Exceptions

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

CLI

PR	Description	Workaround
128503	Oversize frames counter in CLI command 'show interfaces slot/port' is not incrementing when the switch is transmitting/receiving oversize frames.	Use the 'show interfaces slot/port accounting' command and refer to the 'Oversize' parameter.
159817	The keyword 'ssh' is removed from the prompt line after entering 'ssh ? <enter>'	Retype 'ssh' after using the '?'. Retype 'ssh' after using the '?'.

ERP

PR	Description	Workaround
173109	When a ERP-v1 switch is introduced in the ERP-v2 environment the OS-10K CMM shows the version as ERP-v2, however the OS-10K NIs (on which the ERP ports are configured) falls back to V1.	There is no known workaround at this time. This is a display issue only.

Ethernet OAM

PR	Description	Workaround
175885	The RDI value may be set to TRUE for some remote MEPs even though there is no fault in the ethoam domain.	There is no known workaround at this time.

License Management

PR	Description	Workaround
167518	If both Advanced and Datacenter licenses are installed with 7.3.1 on an OS6900 and the switch is downgraded to a previous build, the Advanced license must be reinstalled.	There is no known workaround at this time.

MC-LAG

PR	Description	Workaround
170272	When the MCLAG port goes down on MCLAG-1, the MCLAG-1 chassis still shows the UNP users as local, whereas the MCLAG-2 switch shows the UNP users as remote. Display issue	There is no known workaround at this time.

Qos

PR	Description	Workaround
169730	If the source is not capable of honoring PFC then 5 minutes of continuous PFC frames are seen after 10G port stops oversubscribing.	There is no known workaround at this time
171083	For QoS port ingress rate limiting, IFG plus CRC/preamble (20 bytes) are accounted in addition to the packet size. For QoS port egress rate limiting, these are not accounted. For smaller size packet, the deviation would be higher as compared to the large size packets	There is no known workaround at this time.
172253	On an OS10K when viewing the transmitted and dropped packet count per priority by issuing the command "show qos qsi port 1/1stats", the egress drop statistics will display 0 since the drops occur on the ingress.	There is no known workaround at this time.

SNMP

PR	Description	Workaround
172502	HTTP and HTTPS display enabled on more than one VRF when enabled on the 2nd VRF via SNMP ip service MIB.	Do not use the alaIPService MIB to enable HTTP/HTTPS, use the WebMgt MIB.

Shortest Path Bridging (SPB)

PR	Description	Workaround
172847	When a SPB service is configured an error message similar to "svcNI mSVC error" may appear during an NI hotswap.	There is no known workaround at this time. This is a display issue only, no functional impact.

System

PR	Description	Workaround
175432	With port security enabled on a UNP port, the 'show unip users' command is not showing the complete set of UNP users whereas the MAC table information is correct.	There is no known workaround at this time. This is a display issue only, no functional impact.
176041	Statistics may not be accurate when traffic is being sent from two different NIs and destined for a single NI.	There is no known workaround at this time.

WebView

PR	Description	Workaround
153219	WebView does not display the switch log	View the switch log files from cli using respective commands (more, vi, etc.).
163342	Webview displays more information under the "classification source" that is not available when using the CLI.	There is no known workaround at this time.
171872	When refreshing the Webview SPB SAP statistics page it always starts from the beginning.	To reduce the required presses of the next button to return to a desired row, please use the "Show starting from Service ID" feature. Enter a Service ID, press "Apply" and subsequent page refreshes should start from that Service ID and greatly reduce use of the next button to return to a row.
172391	NTP service on a VRF cannot be enabled/disabled from IP Service webpage	NTP can be enabled/disabled via NTP client/broadcast client mode in Webview or from the CLI.
172493	When using an Internet Explorer 8 in WebView while a page is loading and moving the mouse over a submenu an error message is displayed.	Wait for the page to fully load before moving the mouse over the menu or submenu.

Virtual Chassis

PR	Description	Workaround
173792	NTP updates are not being exchanged between switches when server is in Virtual Chassis mode and client is in standalone mode.	Use the 'ntp src-ip preferred' command to set the EMP-VC IP address to transmit/receive NTP traffic on.
175227	AAA uses EMP-CHAS1 IP address as source IP in VC mode to access LDAP server instead of the EMP-VC IP address.	There is no known workaround at this time.
175482	User cannot authenticate to the slave EMP-IP address in a Virtual Chassis setup using TACACS+, RADIUS, or LDAP.	Authenticate using the local database on the switch.
175771	The command 'reload chassis-id 2 secondary'	There is no known workaround at this time.

	reloads the chassis secondary CMM twice.	
174471	Takeover on master from cmmA to cmmB displays NI down for 3 minutes.	There is no known workaround at this time. This is a display issue only.
176986	With the VC-EMP IP address configured, UNP RADIUS authentication fails due to the RADIUS request packets being sent with the the IP address of the VC-EMP.	Remove the VC-EMP IP address. Note: Removing the VC EMP IP address will require any remote access to be done via the Master Chassis EMP IP address.
177111	When configuring PFC over a VFL on an OS6900, the maximum number of VFL member ports that can be included is 8.	There is no known workaround at this time.

VRF

PR	Description	Workaround
172502	HTTP and HTTPS display enabled on more than one VRF when enabled on the 2 nd VRF via SNMP IP Service MIB.	Do not use the alaIPService MIB to enable HTTP/HTTPS, use the WebMgt MIB.
175622	After deleting and then creating a new VRF it could take up to 3 minutes for traffic to be passed on new VRF.	Perform these steps in order for configurations with the VRF: <ol style="list-style-type: none"> 1. Delete and disable any BGP configuration 2. Remove IP interfaces 3. Delete the VRF and wait for deletion confirmation message

Hot Swap/Redundancy Feature Guidelines

Hot Swap Feature Guidelines

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

Hot Swap Procedure

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

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

```
ChassisSupervisor niMgr info message:  
+++ Expansion module 2 ready!
```

6. Re-insert all transceivers into new module.
7. Re-connect all cables to transceivers.

Hot Swap Time Guidelines

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

Technical Support

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

Region	Phone Number
North America	800-995-2696
Latin America	877-919-9526
Europe	+33-38-855-6929
Asia Pacific	+65 6240 8484

Customers and business partners with an Alcatel-Lucent service agreement may open problem cases 24 hours a day via the Internet.

Worldwide email address: esd.support@alcatel-lucent.com

Worldwide (except North America) web: <https://businessportal.alcatel-lucent.com>

North American Customers: <https://service.esd.alcatel-lucent.com/>

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**.

Also, if needed, we provide all FOSS (Free and Open Source Software) source code used into this release at the following URL: <https://service.esd.alcatel-lucent.com/portal/page/portal/EService/release>

Upgrading an OmniSwitch to 7.3.1.R01

Overview

These instructions document how to upgrade the following OmniSwitch products to 7.3.1.R01 software and firmware. The upgrade must be performed using the CLI. Release 7.3.1.R01 is supported on the following switches. Click on the link to go to instructions for a specific switch.

- [OmniSwitch 10K](#)
- [OmniSwitch 6900](#)

Note: Release 7.3.1.R01 is only supported on **OS10K** and **OS6900** switches. Switches can only be upgraded to Release 7.3.1.R01 using the CLI. Upgrading with WebView is not supported in this release.

Prerequisites

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

- be the responsible party for maintaining the switch's configuration
- be aware of any issues that may arise from a network outage caused by improperly loading this code
- understand that the switch must be rebooted and network users will be affected by this procedure
- have a working knowledge of the switch to configure it to accept an FTP connection through the EMP or Network Interface (NI) Ethernet port
- Read the 7.3.1.R01 GA Release Notes prior to performing any upgrade for information specific to this release.
- Ensure there is a current certified configuration on the switch so that the upgrade can be rolled-back if required.
- Verify the current versions of Uboot on the OS10K. If they meet the minimum requirements, (i.e. they were already upgraded during the 7.2.1.R02 upgrade) then only an upgrade of the AOS images is required.
- Verify the existence of any new modules that were released with 7.3.1.R01 and ensure they are not updated with with an unsupported Uboot version.
- The OS10K must be upgraded to AOS 7.3.1.R01 before inserting any new modules that require AOS 7.3.1.R01.
- Verify the current versions of Uboot and FPGA on the OS6900. If they meet the minimum requirements, (i.e. they were already upgraded during the 7.2.1.R02 upgrade) then only an upgrade of the AOS images is required.

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

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

UBoot, FPGA, Upgrade Requirements

The software versions listed below are the minimum required to run 7.3.1.R01 AOS software.

OmniSwitch 10K

Release	UBoot CMM	UBoot NI	FPGA/CPLD CMM	FPGA/CPLD NI
7.3.1.R01	7.2.1.266.R02	7.2.1.266.R02	No Upgrade Required	No Upgrade Required

Note: Verify the current CMM and NI Uboot versions to determine if a Uboot update is required. New NIs released in 7.3.1.R01 will be shipped with the proper Uboot and do not to be upgraded or downgraded.

OmniSwitch 6900

Release	UBoot CMM	UBoot Expansion Module	FPGA/CPLD CMM	FPGA/CPLD Expansion Module
7.3.1.R01	7.2.1.266.R02	Upgrade Not Supported	1.3.0 1.2.0	No Upgrade Required

Note: Verify the current CMM Uboot and FPGA versions to determine if an update is required.

OmniSwitch 10K – Upgrade Instructions

Upgrading OS10K Switches to 7.3.1.R01 consists of the following steps. The steps should be performed in order:

1. Downloading the Upgrade Files.
2. FTPing the Upgrade Files to the Switch
3. Upgrading the U-Boot File (CMM and NI if required)
4. Upgrading the Image Files.

Note: Upgrading is faster if done in the order above because only one reboot of the switch is required.

Downloading the Upgrade Files

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

- U-Boot File - **u-boot.7.2.1.R02.266.tar.gz**
- Image Files - **Renl.img, Ros.img (7.3.1.R01)**

FTPing the Upgrade Files to the Switch

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

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

Upgrading the U-Boot File

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

CMM Upgrade

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

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

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

Sample output for "update uboot cmm 1"

```
u-boot.bin
u-boot.bin.md5sum
u-boot.bin: OK
[ 482.456298] ffe00000.nor: block unlock error: (status timeout)
Erasing blocks: 4/4 (100%).ease wait.
Writing data: 0k/512k (100%)
Verifying data: 0k/512k (100%)
U-boot successfully updated
Update successfully completed
```

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

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

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

Sample output for "update uboot cmm 2"

```
Please wait...Update successfully completed
```

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

NI Upgrade

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

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

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

Sample output for "update uboot ni all"

```
Please wait....Update successfully completed
```

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

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

IMPORTANT NOTE: Depending on the version of the 7.1.1.R01 Build you are upgrading from, you may receive an error message when you execute the “update uboot ni” command. Simply re-enter the command, and the upgrade will proceed normally (shown below).

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

```
OS10K-> update uboot ni all file u-boot.7.2.1.R02.266.tar.gz
Please wait....Update successfully completed
```

Upgrading the Image Files

Follow the steps below to upgrade the Image Files to 7.3.1.R01:

1. Reload the switch from the working directory.
2. After the switch finishes rebooting, log into the switch.
3. Copy the Image Files from the Working Directory to the Certified Directory.

If you have a **single CMM** switch enter:

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

If you have **redundant CMMs** enter:

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

Sample output for "copy running certified flash-synchro"

```
Wed Feb 8 11:15:35 : flashManager FlashMgr Main info message:
+++ Verifying image directory working on CMM flash
Please wait.Chassis Supervision: CMM has reached the ready state
Chassis Supervision: CMM has reached the ready state
```

WARNING: DO NOT INTERRUPT the upgrade process until it is complete (“CMM has reached the ready state”). Interruption of the process will result in an unrecoverable failure condition.

The upgrade is now complete. See “Verifying the Upgrade” below for information on verifying the upgrade.

Verifying the U-Boot Upgrade

To verify that the U-Boot File was successfully upgraded on the CMM(s), use the show hardware-info command as shown below.

```
OS10K-> show hardware-info

CPU Manufacture           : Freescale Semiconductor
CPU Model                 : MPC 8572
Compact Flash Manufacturer : CF 2GB
Compact Flash size       : 2097930240 bytes
RAM Manufacturer         : Other
RAM size                  : 3998816 kB
CPM FPGA version         : 2.0
U-Boot Version          : 7.2.1.266.R02
CFMs Present              : 1,2,3,4
Power Supplies Present    : 1,2,3,4,-,-,-
Fan Trays Present        : 1,2
NIs Present               : 1,2,3,4,5,6,-,8
```

To verify that the U-Boot File was successfully upgraded on the NI(s), use the show slot command as shown below.

```
OS10K-> show slot 1

Module in slot 1
  Model Name:           OS10K-GNI-C48,
  Description:         10-1000 RJ45,
  Part Number:         907706-90,
  Hardware Revision:   A11,
  Serial Number:       L0360259,
  Manufacture Date:    Nov 10 2011,
  FPGA - Physical 1:   0.7,
  Daughter FPGA - Physical 1: ,
  Daughter FPGA - Physical 2: ,
  Admin Status:        POWER ON,
  Operational Status:  UP,
  Max Power:           152,
  CPU Model Type:      MPC 8572,
  MAC Address:         00:e0:b1:dd:9c:ed,
  ASIC - Physical 1:   BCM56634_B0,
  ASIC - Physical 2:   BCM88230_A0,
  UBOOT Version:    7.2.1.266.R02
```

Verifying the Software Upgrade

To verify that the software was successfully upgraded to 7.3.1.R01, use the show microcode command as shown below.

```
OS10K-> show microcode
```


Package	Release	Size	Description
Ros.img	7.3.1.632.R01	67784336	Alcatel-Lucent OS
Reni.img	7.3.1.632.R01	59189856	Alcatel-Lucent NI

Downgrading Software

Hardware is only backward compatible down to the software that originally supported it. For example, the lowest software version that you can run on an OS10K Switch is 7.1.1.R01. Any new modules released with AOS 7.3.1.R01 cannot be run on any prior version of code. For more information, contact Customer Support.

OmniSwitch 6900 – Upgrade Instructions

Upgrading OS6900 Switches to 7.3.1.R01 consists of the following steps. The steps should be performed in order:

1. Downloading the Upgrade Files.
2. FTPing the Upgrade Files to the Switch
3. Upgrading the U-Boot File. (If required)
4. Upgrading the FPGA. (If required)
5. Upgrading the Image File.

Note: Upgrading in the order above requires only a single reboot of the switch.

Downloading the Upgrade Files

Go to the Alcatel-Lucent Service and Support Website and download and unzip the 7.3.1.R01 upgrade files for the OS6900. The Zip File contains the following files:

- U-Boot File - **u-boot.7.2.1.R02.266.tar.gz**
- FPGA File - **tor_fpgas_130_120.vme**
- Image Files - **Tos.img (7.3.1R01)**

FTPing the Upgrade Files to the Switch

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

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

Upgrading the U-Boot File

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

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

```
OS6900-> update uboot cmm 1 file u-boot.7.2.1.R02.266.tar.gz
```

Sample output for "update uboot cmm 1"

```
u-boot.bin
u-boot.bin.md5sum
u-boot.bin: OK
Erasing blocks: 4/4 (100%)lease wait.
Writing data: 0k/512k (100%)
Verifying data: 0k/512k (100%)

U-boot successfully updated
Update successfully completed
```

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

Upgrading the FPGA

Follow the steps below to upgrade the FPGA File:

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

```
OS6900-> update fpga cmm 1 file tor_fpgas_130_120.vme
```

Sample output for "update fpga cmm 1"

```
Wed Feb  8 11:27:59 : ChassisSupervisor MipMgr info message:
+++ Starting CMM FPGA Upgrade
OS6900 system and expansion fpga update
Please wait.....Update successfully completed
```

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

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

```
OS6900-> rm tor_fpgas_130_120.vme
```

Upgrading the Image File

Follow the steps below to upgrade the Image File:

1. Reload the switch from the working directory.

```
OS6900-> reload from working no rollback-timeout
```

2. After the switch finishes rebooting, log into the switch.

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

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

Sample output for "copy running certified"

```
Wed Feb  8 13:10:17 : flashManager FlashMgr Main info message:
+++ Verifying image directory working on CMM flash
Please wait.....
```

```
Wed Feb  8 13:10:25 : flashManager FlashMgr Main info message:
+++ Image file Tos.img differs - copying file
.....
```

```
Wed Feb  8 13:10:47 : ChassisSupervisor MipMgr info message:
+++ Copy running to certified succeeded
```

WARNING: DO NOT INTERRUPT the upgrade process until it is complete ("Copy running to certified succeeded"). Interruption of the process will result in an unrecoverable failure condition.

The upgrade is now complete. See "Verifying the Upgrade" below for information on verifying the upgrade.

Verifying the U-Boot and FPGA Upgrade

To verify that the U-Boot and FPGA Files were successfully upgraded, use the show hardware-info command as shown below.

```
OS6900-> show hardware-info
```

```

CPU Manufacture           : Freescale Semiconductor
CPU Model                 : MPC 8572
Compact Flash Manufacturer : CF 2GB
Compact Flash size       : 2097930240 bytes
RAM Manufacturer         : Other
RAM size                  : 2071912 kB
FPGA 1 version          : 1.3.0
FPGA 2 version          : 1.2.0
U-Boot Version         : 7.2.1.266.R02
Power Supplies Present   : 1
NIs Present               : 1,2

```

Verifying the Software Upgrade

To verify that the software was successfully upgraded to 7.2.1.R02, use the show microcode command as shown below:

```

OS6900-> show microcode

```

Package	Release	Size	Description
Tos.img	7.3.1.632.R01	106031376	Alcatel-Lucent OS

Downgrading Software

Hardware is only backward compatible down to the software that originally supported it. For example, the lowest software version that you can run on an OS6900 Switch is 7.2.1.R01.

7.3.1.R01 MC-LAG Upgrade Overview


This procedure is documented in order to help reduce the overall downtime when upgrading to AOS Release 7.3.1.R01 in an MC-LAG environment. A Layer 3 OSPF topology is used as an example; and while not every topology will be exactly the same, most L3 redundant topologies will be similar to the diagram below. The topology below uses OS10Ks, however, the upgrade procedure also applies to an MC-LAG topology using OS6900s.

Due to continuous enhancements to the MC-LAG feature throughout AOS Release 7, it's a requirement that Multi-Chassis Peer switches be running the same AOS release version. To meet this requirement with minimal disruption the following key points should be followed:

- Each Multi-Chassis Peer has to be rebooted one at a time to complete the upgrade requirement with minimal disruption.
- While performing the upgrade the time that the two peer switches are running different versions of code is minimized.
- Flushing of the L3 tables at the proper time will ensure that the transition and synchronization of the L3 databases between the peer switches is handled properly between the different versions of code.

Following the steps documented below will ensure minimal traffic disruption during the upgrade procedure. While the time will vary based on the topology and chassis configuration, a successful upgrade was performed in Alcatel-Lucent's SQA lab with just 3 seconds of disruption.

Note: This upgrade procedure is only supported when upgrading from AOS Release 7.2.1.R02 to 7.3.1.R01. This is due to MC-LAG enhancements such as Chassis Group ID and additional platform feature parity incorporated into 7.2.1.R02 which are not in earlier AOS Release 7 versions.



Topology Example

7.3.1.R01 MC-LAG Upgrade Prerequisites

Based on the example topology above, Core1 and Core2 are the Multi-Chassis Peer switches. Multi-Chassis Peer switches **MUST** be running the same AOS release version. In order to meet the requirement, both switches must be rebooted one at a time to complete the 7.3.1.R01 AOS upgrade. The following prerequisites should be understood prior to performing the upgrade:

1. This procedure is only supported on peer switches that are running AOS Release 7.2.1.R02.
2. Rebooting each peer switch will result in downtime of that switch from between 5 to 10 minutes depending on the configuration. However, during this downtime traffic will be re-routed to the other peer switch.
3. Any devices directly connected to the peer switch being rebooted will not be reachable while the switch is rebooting.
4. Read and understand the 7.3.1.R01 Upgrade Instructions section in the AOS 7.3.1.R01 Release Notes for AOS and Uboot requirements.

Upgrade Procedure

In order to minimize the downtime for the upgrade, the following steps are recommend based on the topology example given above.

Prepare for the Upgrade

1. Download the necessary image files from Service & Support as described in the 7.3.1R01 Upgrade Instructions.
2. Open a direct console connection to both Core1 and Core2.
3. Save the configuration and perform a flash synchronization on both peer switches by issuing the following command:

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

4. FTP the 7.3.1.R01 GA AOS images to both Core1 and Core2

Core1 Upgrade

1. Reload Core1 by issuing the following command:

```
Core1-> reload from working no rollback-timeout
```

2. Wait for Core1 to reboot and become operational. Depending on the configuration this may take from 5 to 10 minutes. During the reboot of Core1 traffic will be redirected through Core2.
3. Via the console connection to Core1, check the module status and wait for OSPF adjacency to be established.
4. During Core1's OSPF adjacency / MC-LAG establishing – there is a SPLIT chassis between Core1 (new code) and Core2 (old code) which may result in ARP and L3 routing table inconsistencies. In order to address these inconsistencies, the ARP table will be flushed prior to rebooting Core2.

Core2 Upgrade

1. Disable all links on Core2 (including VFL ports, OSPF links and MC-LAG links) to prevent a SPLIT chassis scenario when Core2 is rebooted:

```
Core2-> interfaces <slot> admin-state disable
```

2. Clear the ARP table on Core1 to remove any L3 inconsistencies that may have been created due to Core1 having a different version of code than Core2:

```
Core1-> clear arp-cache
```

3. Verify that all traffic has recovered prior to reloading Core2, then enter the following:

```
Core2-> reload from working no rollback-timeout
```

4. Once Core2 reboots and Multi-Chassis is once again operational, there will be one more brief re-convergence.