Release Notes

OmniSwitch 6860/6860E

Release 8.1.1.R01

These release notes accompany release 8.1.1.R01 software which is supported on the OmniSwitch 6860/6860E 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.

Contents

Related Documentation

These release notes should be used in conjunction with OmniSwitch AOS Release 8 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/UserGuides

OmniSwitch 6860/6860E Hardware User Guide
Describes the hardware and software procedures for getting an OmniSwitch up and running as well as complete technical specifications and procedures for all OmniSwitch chassis, power supplies, and PoE.

OmniSwitch AOS Release 8 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 8 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 8 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, Layer 3 information, security options, and Quality of Service.

OmniSwitch AOS Release 8 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, IS-IS, OSPF, and OSPFv3.

OmniSwitch AOS Release 8 Transceivers Guide
Includes 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 6860/6860E Series Release 8.1.1.R01 requires 2GB of RAM 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

OmniSwitch 6860/6860E models released in 8.1.1.R1 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.

OmniSwitch 6860/6860E (All models) – AOS Release 8.1.1.497.R01(GA)

Model	Uboot	FPGA
OS6860/OS6860E (except U28)	8.1.1.70.R01	Version 0.9
OS6860E-U28	8.1.1.70.R01	Version 1.4

New Hardware Support in 8.1.1.R01

OmniSwitch 6860-24 - Fixed configuration chassis in a 1U form factor with 24 10/100/1000 BaseT ports, 4 fixed SFP+ (1G/10G) ports and two 20G virtual fabric link ports.

OmniSwitch 6860-48 - Fixed configuration chassis in a 1U form factor with 48 10/100/1000 BaseT ports, 4 fixed SFP+ (1G/10G) ports and two 20G virtual chassis link ports.

OmniSwitch 6860-P24 - Fixed configuration chassis in a 1U form factor with 24 10/100/1000 BaseT PoE ports, 4 fixed SFP+ (1G/10G) ports and two 20G virtual chassis link ports.

OmniSwitch 6860-P48 - Fixed configuration chassis in a 1U form factor with 48 10/100/1000 BaseT PoE ports, 4 fixed SFP+ (1G/10G) ports and two 20G virtual chassis link ports.

OmniSwitch 6860E-24 - Fixed configuration chassis in a 1U form factor with 24 10/100/1000 BaseT ports, 4 fixed SFP+ (1G/10G) ports, two 20G virtual chassis link ports, and an EMP port for out-of-band management. Includes a built-in co-processor for enhanced network services.

OmniSwitch 6860E-48 - Fixed configuration chassis in a 1U form factor with 48 10/100/1000 BaseT ports, 4 fixed SFP+ (1G/10G) ports, two 20G virtual chassis link ports, and an EMP port for out-of-band management. Includes a built-in co-processor for enhanced network services.

OmniSwitch 6860E-P24 - Fixed configuration chassis in a 1U form factor with 24 10/100/1000 BaseT PoE ports, 4 fixed SFP+ (1G/10G) ports, two 20G virtual chassis link ports, and an EMP port for out-of-band management s. Includes a built-in co-processor for enhanced network services. Supports 60W of PoE on ports 1-4 and 30W on all other ports.

OmniSwitch 6860E-P48 - Fixed configuration chassis in a 1U form factor with 48 10/100/1000 BaseT PoE ports, 4 fixed SFP+ (1G/10G) ports, two 20G virtual chassis link ports, and an EMP port for out-of-band management. Includes a built-in co-processor for enhanced network services. Supports 60W of PoE on ports 1-4 and 30W on all other ports.

OmniSwitch 6860E-U28 - Fixed configuration chassis in a 1U form factor with 28 ports supporting 1000BASE-X and 100BASE-FX, 4 fixed SFP+ (1G/10G) ports, two 20G virtual chassis link ports, and an EMP port for out-of-band management. Includes a built-in co-processor for enhanced network services.

OmniSwitch AC System Power Supply (OS6860-BP/PS-150-AC) - Modular AC power supply. Provides 150W system power to one OS6860 non-PoE switch. **Note**: Mixing of different types of wattage power supplies in the same chassis is not supported.

OmniSwitch DC System Power Supply (OS6860-BP-D/PS-150-DC) - Modular DC power supply. Provides 150W system power to one OS6860 non-PoE switch. **Note**: Mixing of different types or wattage power supplies in the same chassis is not supported.

OmniSwitch AC PoE Power Supply (OS6860-BP-PH/PS-600W-AC-P) - Modular 600W AC PoE power supply. Provides system and 450W of PoE power to one 24 port PoE switch. **Note**: Mixing of different types or wattage power supplies in the same chassis is not supported.

OmniSwitch AC PoE System Power Supply (OS6860-BP-PX/PS-920W-AC-P) - Modular 920W AC PoE power supply. Provides system and 780W of PoE power to one 48 port PoE switch. **Note**: Mixing of different types or wattage power supplies in the same chassis is not supported.

OmniSwitch non-PoE Fantray – Modular fantray designed to be used with non-PoE models using the OS-BPS as as backup power supply. Provides chassis cooling in case of the loss of the main power supply.

OmniSwitch Backup Power Shelf (OS-BPS) - The OmniSwitch Backup Power Shelf/System (BPS) is a 1RU power shelf that provides a backup system power to up to 8 OS6860/6860E non-PoE switches. A maximum of two system power supplies can be installed in a single OS-BPS chassis.

- Supports up to eight switches
- Supports up to 900W of redundant system power

OmniSwitch BPS System Power Supply (OS-BPS-S) - 450W System Power Supply for the OS-BPS (OS-PS-450W-A).

- Provides 450W of redundant system power
- The OS-BPS can support 2 OS-BPS-S power supplies

Transceivers

Gigabit	Dual Speed	100M	10-Gigabit	VFL stacking cables
SFP-GIG-SX*	SFP-DUAL-MM	SFP-100-LC-MM	SFP-10G-SR	QSFP-40G-SR*
SFP-GIG-LX	SFP-DUAL-SM10*	SFP-100-LC-SM15	SFP-10G-LR	OS6860-CBL-100 (1m)
SFP-GIG-LH40	SFP-DUAL-BX-D	SFP-100-LC-SM40	SFP-10G-ER	OS6860-CBL-300 (3m)
SFP-GIG-LH70	SFP-DUAL-BX-U	SFP-100-BXLC-D*	SFP-10G-LRM	OS6860-CBL-40 (40cm)
SFP-GIG-EXTND		SFP-100-BXLC-U*	SFP-10G-GIG-SR	
SFP-GIG-T			SFP-10G-C	
SFP-GIG-BX-D				
SFP-GIG-BX-U				
SFP-GIG-BX-D20				
SFP-GIG-BX-U20				
SFP-GIG-BX-D40				
SFP-GIG-BX-U40				
SFP-GIG-CWD				

* Not supported on the OS6860E-U28. Check for availability.

New Features Summary

The following software features are being introduced with the 8.1.1.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' require the installation of that license.

AOS 8.1.1.R01 Feature Summary Table

Feature	Platform	License
Hardware / Virtual Chassis Feature Support		
micro-USB, RS-232, and Bluetooth console access	OS6860/6860E	Base
IEEE 802.3ah Dying Gasp	OS6860/6860E	Base
IEEE 802.3az (EEE)	OS6860/6860E	Base
OmniSwitch Backup Power Shelf (OS-BPS)	OS6860/6860E	Base
Power over Ethernet (PoE)	OS6860/6860E	Base
USB Support	OS6860/6860E	Base
Virtual Chassis (VC)	OS6860/6860E	Base
Virtual Chassis Split Protection (VCSP)	OS6860/6860E	Base
Manageability Feature Support		
CLI and CLI Abbreviation	OS6860/6860E	Base
Ethernet Interfaces	OS6860/6860E	Base
In-Service Software Upgrade (ISSU)	OS6860/6860E	Base
License Management	OS6860/6860E	Base
LLDP Network / Voice VLAN Policies	OS6860/6860E	Base
Multiple VRF Routing and Forwarding	OS6860/6860E	Advanced
Network Time Protocol (NTP)	OS6860/6860E	Base
OpenFlow Agent	OS6860/6860E	Base
Partitioned Switch Management	OS6860/6860E	Base
Pause Control/Flow Control	OS6860/6860E	Base
Remote Access • FTP • SCP • SSH/SFTP • Telnet • TFTP	OS6860/6860E	Base

Feature	Platform	License
Remote Configuration Download	OS6860/6860E	Base
SNMP	OS6860/6860E	Base
Software Rollback – Multi-Image/Multi-Config	OS6860/6860E	Base
Storm Control	OS6860/6860E	Base
Text File Configuration	OS6860/6860E	Base
Web-Based Management (WebView)	OS6860/6860E	Base
Web Services & CLI Scripting	OS6860/6860E	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS6860/6860E	Base
802.1Q	OS6860/6860E	Base
IEEE 802.1q Shortest Path Briding (SPB)	OS6860/6860E	Advanced
Automatic VLAN Containment (AVC)	OS6860/6860E	Base
Ethernet Ring Protection v2 (ERPv2)	OS6860/6860E	Base
High Availability VLAN	OS6860/6860E	Base
Link Aggregation –Static and LACP (802.3ad/ax) - Non-unicast load balancing	OS6860/6860E	Base
Multiple VLAN Registration Protocol (MVRP)	OS6860/6860E	Base
Source Learning - Disable learning on a port - Disable learning on a VLAN	OS6860/6860E	Base
Spanning Tree - 802.1d and 802.1w - Multiple Spanning Tree Protocol - PVST+	OS6860/6860E	Base
Layer 3 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS6860/6860E	Base
Bind physical port to an IP interface	OS6860/6860E	Base
Border Gateway Protocol (BGP) 4 with Graceful Restart	OS6860/6860E	Advanced

Feature	Platform	License
Border Gateway Protocol (BGP) 4 with 4-octet ASN Support	OS6860/6860E	Advanced
DHCP / UDP - DHCP Relay/Option-82 - Per-VLAN DHCP Relay - UDP Relay	OS6860/6860E	Base
DNS Client	OS6860/6860E	Base
Generic Routing Encapsulation (GRE)	OS6860/6860E	Base
IP Multinetting	OS6860/6860E	Base
IP Route Map Redistribution	OS6860/6860E	Base
IP-IP Tunneling	OS6860/6860E	Base
Intermediate System to Intermediate System (IS-IS) for IPv4	OS6860/6860E	Advanced
Non-Contiguous Mask and IPv6 Gateway Support	OS6860/6860E	Base
Open Shortest Path First (OSPF) v2	OS6860/6860E	Advanced
OSPF – convert interfaces into OSPF passive interfaces	-	Advanced
Routing Information Protocol (RIP) v1/v2	OS6860/6860E	Base
Recursive Static Route	OS6860/6860E	Base
Routing to an IP interface name	OS6860/6860E	Base
Routing Protocol Preference	OS6860/6860E	Base
Server Load Balancing (SLB) - SLB WRR	OS6860/6860E	Base
Session Initiation Protocol (SIP) Snooping	OS6860/6860E	Base
Virtual Router Redundancy Protocol (VRRP) v2	OS6860/6860E	Base
IPv6 Feature Support		
- ECMP routes for IPv6 - IPv6 Routing - IPv6 Client and/or Server Support - Globally Unique Local Unicast Addresses - IPv6 Scoped Multicast Addresses - IPv6 Neighbor Discovery	OS6860/6860E	Base

Feature	Platform	License
Border Gateway Protocol (BGP) 4 and BGP IPv6 Extensions	OS6860/6860E	Advanced
IPv6 DHCP Relay	OS6860/6860E	Base
IPSec - IPv6 - OSPFv3 - RIPng	OS6860/6860E	Advanced
Intermediate System to Intermediate System (IS-IS) IPv6	OS6860/6860E	Advanced
Open Shortest Path First (OSPF) v3 - point to point OSPFv3 interfaces		Advanced
RIPng	OS6860/6860E	Advanced
SIP Snooping	OS6860/6860E	Base
VRRPv3	OS6860/6860E	Base
QoS Feature Support		
Auto-Qos Prioritization of NMS Traffic	OS6860/6860E	Base
Auto-Qos Prioritization of IP Phone Traffic	OS6860/6860E	Base
Ingress and egress bandwith shaping	OS6860/6860E	Base
Policy Based Routing (Permanent Mode)	OS6860/6860E	Advanced
Redirect Policies (Port and Link Aggregate)	OS6860/6860E	Base
Tri-Color Marking	OS6860/6860E	Base
VFC/VoQ Profiles - Four pre-defined Qset Profiles - WRED is not supported	OS6860/6860E	Base
Multicast Feature Support		
IP Multicast Switching (IGMP)	OS6860/6860E	Base
IP Multicast Switching (Proxying)	OS6860/6860E	Base
DVMRP	OS6860/6860E	Advanced
IGMP Multicast Group Configuration Limit	OS6860/6860E	Base
IGMP Relay	OS6860/6860E	Base
L2 Static Multicast Address	OS6860/6860E	Base
PIM-SM/PIM-DM/PIM-SSM (Source-Specific Multicast)	OS6860/6860E	Advanced

Feature	Platform	License
PIM/DVMRP Interoperability	OS6860/6860E	Advanced
Monitoring/Troubleshooting Feature Support		
Digital Diagnostic Monitoring (DDM)	OS6860/6860E	Base
Fault Propagation and Link Flapping - Wait to Shutdown - Wait to Restore	OS6860/6860E	Base
Gigaword Packet Counters	OS6860/6860E	Base
Health Statistics	OS6860/6860E	Base
Line Diags & Enhanced Port Performance (EPP)	OS6860/6860E	Base
Ping and Traceroute / Extended Ping and Traceroute	OS6860/6860E	Base
Port Mirroring - Policy Based Mirroring - Remote Port Mirroring	OS6860/6860E	Base
Port Monitoring	OS6860/6860E	Base
Rmon	OS6860/6860E	Base
sFlow	OS6860/6860E	Base
Switch Logging and Syslog	OS6860/6860E	Base
Time Domain Reflectometry (TDR)	OS6860/6860E	Base
Uni-Directional Link Detection (UDLD)	OS6860/6860E	Base
Metro Ethernet Feature Support		
Ethernet OAM / ITU Y1731 and 802.1ag	OS6860/6860E	Base
ERP G.8032 – Shared VLAN	OS6860/6860E	Base
Ethernet Services	OS6860/6860E	Base
L2 Control Protocol Tunneling (L2CP)	OS6860/6860E	Base
Security Feature Support		
Application Monitoring (Appmon)	OS6860/6860E	Base
Access Control Lists (ACLs) for IPv4/IPv6	OS6860/6860E	Base

Feature	Platform	License
Access Guardian 2.0	OS6860/6860E	Base
Account and Password Policies	OS6860/6860E	Base
Admin User Remote Access Control	OS6860/6860E	Base
ARP Defense Optimization	OS6860/6860E	Base
ARP Poisoning Detect	OS6860/6860E	Base
Authenticated Switch Access	OS6860/6860E	Base
IP DoS Filtering	OS6860/6860E	Base
Learned Port Security (LPS)	OS6860/6860E	Base
Policy Server Management	OS6860/6860E	Base
Port Mapping (Private VLANs)	OS6860/6860E	Base

Early Availability / Demonstration Features Summary

The following software features are being introduced with the 8.1.1.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.

AOS 8.1.1.R01 Early Availability / Demonstration Summary Table

Feature	Platform	License
Deep Packet Inspection (DPI)	OS6860/6860E	Base

New Features Descriptions

Hardware/Virtual Chassis Features

USB Console, RS-232, and Bluetooth console access

Access to the OS6860 console can be accomplished using the various methods listed below. A micro-USB to USB console cable is shipped with each OS6860, a USB to UART driver is required to use this cable for console access. Refer to the OmniSwitch AOS Release 8 Switch Management Guide for additional information.

- **USB Console –** Cable included with shipment. A USB to UART driver must be installed on the laptop/device connecting to the OmniSwitch. Driver and installation instructions can be downloaded from: http://www.silabs.com/products/mcu/pages/usbtouartbridgevcpdrivers.aspx
- **RS-232 –** A micro-USB to RJ-45 adapter is required. Check for availability.
- **Bluetooth –** A bluetooth USB adapter can be connected to the USB port and used to provide console access. The following bluetooth devices were validated:
  - TRENDnet TBW-107UB - Network adapter - USB - Bluetooth 2.1 EDR - Class 2
  - ZOOM 4314 USB Adapter - Network adapter - USB - Bluetooth 2.1 EDR - Class 1
  - Belkin USB 4.0 Bluetooth Adapter - Network adapter - USB - Bluetooth 4.0
  - IOGEAR Bluetooth 4.0 USB Micro Adapter Multi-Language Version - Network adapter - USB - Bluetooth 4.0 - Class 2
  - SMK-Link Electronics Nano Bluetooth Dongle 4.0 LE + EDR - Network adapter - USB 2.0 - Bluetooth 4.0 EDR)

IEEE 802.3ah Dying Gasp

This feature is designed to send a message on power loss. There are three types of messages sent:

SNMP Trap

As soon as the power failure is detected, a SNMP trap message is sent to the first three configured SNMP stations. The trap includes the following information:

- Slot number
- Power supply type (primary/backup)
- Time of the failure

Syslog Message

As soon as the power failure is detected a syslog message is sent to the first four syslog servers configured.

Link OAM PDU

As soon as the power failure is detected, an 802.3ah OAM Information PDU is sent to all ports of the NI for which link OAM is enabled. The PDU will have the Dying Gasp bit set.

IEEE 802.3az (EEE)

Energy Efficient Ethernet (EEE) is a protocol to allow ports to operate in idle or low power mode when there is no traffic to send. When EEE is enabled on a port it will advertise its EEE capability to its link partner. If the partner supports EEE they will operate in EEE mode. If the partner does not support EEE the ports will operate

in legacy mode. This allows EEE capable switches to be deployed in existing networks avoiding backward compatibility issues.

- EEE is only applicable to OmniSwitch copper ports operating at 100/1000 Mbps speed.
- The LLDP option in IEEE 802.3az standard is not currently supported.

OmniSwitch Backup Power Shelf (OS-BPS)

The OmniSwitch BPS is a 1.5U power shelf to support a backup power solution for the OmniSwitch 6860/6860E non-PoE switches. The OS-BPS can be used to provide backup power to a stack of 8 OmniSwitches or 8 standalone OmniSwitches and help to reduce the required rack space. The amount of power supplied by the OS-BPS depends on the number of system PoE power supplies installed as noted in the table below:

- 1 System P/S - Can support up to 450W of backup system power.
- 2 System P/S - Can support up to 900W of backup system power.

How the OS-BPS handles the failure of a primary power supply installed in an OmniSwitch depends on the mode the OS-BPS is operating in as described below.

Note: OmniSwitch 6850s and OmniSwitch 6860s are not supported on the same OS-BPS. The OS-BPS is not supported with OS6860E PoE models. The OS6860(E) non-Poe models can have both chassis power supplies or a chassis power supply and fantray when used with an OS-BPS.

Single Mode (N+1)

When configured in single mode the OS-BPS power is unmanaged and does not provide load sharing. The OS-BPS will only provide backup power if all the internal power supplies are no longer functioning. There is no priority given to any connector, the OS-BPS will continue to provide power as long as the amount of power required by the OmniSwitches is less than the total power available on the OS-BPS. The number of OmniSwitch primary power supplies that can fail but still have redundant system power supplied by the OS-BPS depends on the number of OS-BPS system power supplies installed.

Full Mode (N+N)

When configured in full mode the OS-BPS power is managed, meaning if the power requirements of the switches becomes greater than what the OS-BPS can provide, the OS-BPS will manage the power and begin to shut down power based on connector priority.

The OS6860(E) non-Poe models can have both chassis power supplies or a chassis power supply and fantray when used with an OS-BPS. See the summary table below.

Power over Ethernet (PoE)

Power over Ethernet (PoE) provides inline power directly from the switch's Ethernet ports. From these RJ-45 ports the devices receive both electrical power and data flow. The switch supports both both IEEE 802.3af and 802.3at standards. All PoE models support 802.3at on all PoE ports, the OS6860E PoE models support up to 60W on the first 4 ports.

Class Detection - Allows the OmniSwitch to provide up to 30W of PoE power as well as automatically detect the Class (Class 0, Class1, Class2, Class3 or Class4) of the connected powered device. This allows the OmniSwitch to automatically adjust the maximum allowed power for a port preventing the OmniSwitch from delievering more power than the device requires.

PoE Power Rules - Allows an administrator to assign specific events or behaviors to PoE ports. For example, a user can set up a rule that powers off PoE devices at a specific time of day, then restores power after a specified amount of time has elapsed. Events can also be configured to take place on specific days of the week, as well as specific months of the year. Power rules are used in conjunction with power policies. Users must first configure power rules on the switch, then assign the rules to a policy. Once the rules have been assigned, the policy (and its accompanying rules) can be bound to specific slots or ports on the switch or Virtual Chassis.

Priority Disconnect - Since not all Powered Devices (PDs) connected to the switch have the same priority within a the network, the OmniSwitch allows the administrator to specify priority levels on a port-by-port basis. Priority levels include low, high, and critical. These values are used by the switch when determining whether an incoming or existing PD will be granted or denied power when there are too few watts remaining in the PoE power budget for all PoE devices. If the PoE power budget cannot support the PoE draw, the system will determine which devices will be powered. Based on priority disconnect rules, in some cases one or more existing devices may be powered down in order to accommodate a higher priority incoming device, or the loss of a power supply.

Note: Priority disconnect is supported for up to 450W on each 600W PoE power supply. Priority disconnect is supported for up to 780W on each 920W PoE power supply.

USB Support

The USB port can be used with an Alcatel-Lucent certified USB Flash drive 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.

Virtual Chassis (VC)
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. Up to eight switches, in any combination of chassis types, can be combined into a single virtual chassis. A VC can easily expand switching capacity simply by adding additional switches to the VC. For example, a deployment can start with a VC composed of two switches with the option of adding up to six additional switches to that VC as network demands increase over time. VCs also provide enhanced resiliency and redundancy features. If a switch in a VC goes down or is taken offline, the other switches in the VC will continue to operate without disruption. In addition operating software and configuration parameters are synchronized on all switches in the VC.

Note: All OS-6860s are considered to be a virtual chassis. A single OS-6860 may be referred as a standalone chassis but still requires the *chassis/slot/port* CLI notation. A standalone OS-6860 should not be confused with "standalone mode" that was introduced in previous versions of AOS that supported the Virtual Chassis feature.

The following are some key points regarding Virtual Chassis configuration:

- All OmniSwitch 6860s run as a Virtual Chassis. A single OS6860 is considered a "virtual-chassis-of-one" but may also be referenced as a standalone chassis.
- A Virtual Chassis can 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.
- Virtual Chassis requires a chassis identifier to be used for all slot and port related command. (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 the dedicated VFL ports.

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).

Note: The Virtual Chassis EMP addresses are only supported on OS680E models (non-E models do not have an EMP port). When configuring VC EMP addresses in a VC with a mix of OS6860E models and OS6860-non-E models the VC could become unreachable via the EMP port if a non-E model becomes Master. In a mixed VC environment it's recommended to create a management VLAN for remote management purposes.

Virtual Chassis Split Protection (VCSP)

In the case of a virtual chassis splitting into disjoint sub-VCs due to the failure of one or more VLFs both of the resulting VCs could end up having the same system MAC and IP addresses. Since there is no communication between these individual VCs due to the VFL failure they end up communicating with the rest of the network devices using the same MAC and IP addresses. This VC split scenario is disruptive to the network as the conflicting MAC and IP addresses can lead to layer 2 loops and layer 3 traffic disruption.

VCSP provides the following benefits:

- Avoid network disruptions by preventing duplicate MAC and IP addresses on the network.
- The sub-VC that forms out of the VC split is able to detect that a split has occurred by use of a helper switch. The helper functionaltiy is supported on an OS6860/OS6860E, OS6850E, OS9000E, or OS6450 (with the appropriate 6.6.4 maintenance release).

- Once the VC split condition has been determined, the sub-VC will put its front-panel ports into an operationally down state preventing traffic forwarding and avoiding loops and possible traffic disruption. The VCSP link aggregate ports will remain up.
- A trap can be sent by the active-VC indicating the VC split state. The trap indicates that the split has occured and which elements are in the operationally down sub-VC.
- The entire VC will automatically recover when the sub-VC rejoins the VC.

This feature can also be leveraged for detecting a VC split in a remote VC topology where the VC may consist of elements located in different physical locations.

Note: A redundant VFL cable should be used for best traffic convergence in the event of failure.

Manageability Feature Support

Command Line Interface (CLI)

The command line interface (CLI) is a Bash-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.

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

Ex. "show vlan" can be abbreviated to "sh vl"

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, inter-frame gap, 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 images running on a virtual chassis with minimal disruption to data traffic. The images can be upgraded on a fully synchronized, certified, and redundant virtual chassis running an ISSU capable build.

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 various advanced routing protocols as listed in the feature summary table.

LLDP Network / Voice VLAN Policies

LLDP Network policy allows the advertisement of VLAN id, 802.1p and DSCP for the following applications:

Voice, Voice Signaling, Guest Voice, Guest Voice Signaling, Soft phone voice, Video Conferencing, Streaming voice and Video Signaling.

The OmniSwitch use LLDP-MED Network Policies to advertise the Voice VLAN to the connected IP Phones through explicit definition of LLDP-MED Network Policy that contains information about the VLAN-ID and the associated L2 and L3 priorities. The binding of the network policies can be done globally or on a per port basis. The VLAN must be created explicitly. When using authenticated or mobile VLANs it is recommended to use mobile-tag rules to dynamically associate the devices according to the incoming tagged traffic

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 can exist 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 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.

- **VRF – PIM and DVMVRP**

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

- **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 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.

VRF Profiles

VRF profiles are used to increase the number of supported VRFs. Two types of profiles can be configured called 'low' and 'max'. A 'max' profile has no restrictions on the number of IPv4 protocols.

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

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

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 software by default will be able to respond to NTP client requests, and establish a client/server peering relationship.

OpenFlow Agent

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

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

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

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

Note: OpenFlow is only supported on a standalone chassis.

Partitioned Switch Management

A user account includes a login name, password, and user privileges. The privileges determine whether the user has read or write access to the switch, and which command domains and command families the user is authorized to execute on the switch. The privileges are sometimes referred to as authorization; the designation of particular command families or domains for user access is sometimes referred to as partitioned management.

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. The OmniSwitch does not support the transmission of PAUSE frames but will honor received PAUSE frames.

Remote Access

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

Secure Copy (SCP)
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 unsecure 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.

Remote Configuration Download

The Remote Configuration Download capability automates and simplifies the deployment of large network installations eliminating the need for manual configuration of each device. It also ensures that each device is compliant with the centrally controlled device configuration policies and firmware revisions. This feature allows a newly deployed OmniSwitch to automate the process through an instruction file that provides the necessary actions to download its configuration or any necessary firmware upgrades with no user intervention by doing the following:

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

DHCP on VLAN 1, VLAN 127, LLDP tagged management VLAN

DHCP client timing out in few minutes causes operational concerns. Hence DHCP client operation will continuously try to obtain a DHCP lease alternating between the following methods:

- Static DHCP client on untagged VLAN 1
- Dynamic DHCP client on tagged VLAN 127
- Dynamic DHCP client on LLDP tagged management VLAN

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.

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

Web Services & CLI Scripting

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

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

Layer 2 Feature Descriptions

802.1AB MED Extensions

IEEE 802.1AB is a standards based connectivity discovery protocol. The purpose of the IEEE standard 802.1AB for Link Layer Discovery Protocol (LLDP), is to provide support for network management software dealing with topology discovery. Switches that are compliant with 802.1AB exchange information with neighboring devices and maintain a database of the information exchanged. 802.1ab uses TLV (Time, Length, Value) frames to exchange information with neighboring devices.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.

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 Provide 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

Automatic VLAN Containment (AVC)

In an 802.1s Multiple Spanning Tree (MST) configuration, it is possible for a port that belongs to a VLAN, which is not a member of an instance, to become the root port for that instance. This can cause a topology change that could lead to a loss of connectivity between VLANs/switches. Enabling Automatic VLAN Containment (AVC) helps to prevent this from happening by making such a port an undesirable choice for the root. When AVC is enabled, it identifies undesirable ports and automatically configures them with an infinite path cost value. Balancing VLANs across links according to their Multiple Spanning Tree Instance (MSTI) grouping is highly recommended to ensure that there is not a loss of connectivity during any possible topology changes. Enabling AVC on the switch is another way to prevent undesirable ports from becoming the root for an MSTI.

Ethernet Ring Protection Version 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 currently supported:

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

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/802.3ax)

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.

Multiple VLAN Registration Protocol (MVRP)

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 opology

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.

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 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 the default Spanning Tree protocol for the switch regardless of which mode is active.

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.

Layer 3 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.

Bind physical port to an IP interface (Routed VLAN)

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

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 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.

BGP 4-Octet Autonomous System Number (ASN)

This feature provides the following:

- BGP Support for 4-octet (32 bit) ASN for BGP neighbor interoperability and path attribute interoperability as per RFC 6793.
- Advertisement and discovery of 4-octet ASN capability by using the BGP Capability advertisement as specified in RFC 5492.
- Support for two new optional transitive attributes AS4_PATH and AS4_AGGREGATE. These attribute are used when new BGP speakers are interacting with OLD BGP speakers.
- To establish a neighbor relationship between non-mappable BGP 4-octet ASNs with BGP 2-octet ASNs the reserved 2-octet ASN AS_TRANS 23456 is used.
- The 4-octet AS Specific Extended Community as specified in RFC 5668 will be used with non-mappable 4-octet ASNs. If the ASN is mappable to 2-octet, the 2-octet AS specific extended community will still be used.
- The 4-octet ASN is represented in one of three ways:
    - asplain (simple decimal notation)
    - asdot+ (two 16-bit values as low-order and high-order)
    - asdot (a mixture of asplain and asdot+)

DHCP / UDP Relay

DHCP Relay allows for forwarding of DHCP broadcast requests to configurable DHCP server IP address in a routing environment. DHCP snooping provides network security by filtering untrusted DHCP messages by

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

Ingress Source Filtering

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

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

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

DHCP Relay Agent Information Option-82

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

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

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

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 (GRE)

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 multiple 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.

Intermediate System to Intermediate System (IS-IS) for 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. This release supports multi-VRF aware IS-IS for IPv4.

Non-contiguous Mask and IPv6 Gateway Support

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

Open Shortest Path First (OSPF) v2

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.

Routing Information Protocol (RIP) v1/v2

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.

Recursive Static Route

Recursive static routes are similar to static routes. However, with a recursive static route the gateway does not have to be a directly connected router. If the OmniSwitch is unable to find a route in the routing table for a packet it can use the recursive static route. The OmniSwitch will use its routing table to lookup a route for the gateway instead of having to use a directly connected router. This feature can be used in large networks to configure a uniform static route for all routers on a network. Each router will use the same gateway but the path to reach the gateway may differ for each router.

Routing to an IP interface name

In case of directly connected NAT routers an interface name can be used instead of gateway IP address, provided the router is enabled for proxy-ARP to handle ARP requests for the route addresses.

Providing the interface name makes it easy to configure static routing through a NAT router. The only information that needs to be provided is the name of the interface that will be used, not a specific IPv4 gateway address on the NAT router.

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.

Session Initiation Protocol (SIP) Snooping

SIP Snooping feature address the key challenge of real time delivery and monitoring requirements for media streams from SIP devices. SIP snooping provides plug and play support to the device, where it automatically identifies the ports used. It also enhances the security of device. SIP Snooping prioritizes voice and video traffic over non-voice traffic. To summarize, SIP Snooping:

- Identifies and marks the SIP and its corresponding media streams. Each media stream contains RTP and RTCP flows. Marking is done using the DSCP field in IP header.
- Provides user configured QOS treatment for SIP/RTP/RTCP traffic flows based on its marking.
- Calculates QOS metric values of delay, jitter, round trip time, R factor and MOS values of media streams from its corresponding RTCP.

Virtual Router Redundancy Protocol (VRRP) v2

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. 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.

Local Proxy Neighbor Discovery (LPND)

Local Proxy Neighbor Discovery (LPND) is used to isolate IPv6 nodes on the same VLAN from each other. If LPND is enabled on an IPv6 VLAN interface, a client will not learn the MAC address of any other IPv6 node reached via the switch. The switch will intercept all neighbor discovery messages and replace the client MACs with the switches MAC before sending the messages to their destination. As a result, all IPv6 traffic will be routed not switched.

Router Advertisement (RA) Filtering

RA filtering can be used to prevent the spread of rogue RAs from unauthorized systems. If enabled on an interface, any received RAs will be dropped without being forwarded on to any other connected IPv6 clients. One or more trusted ports or linkaggs can be specified for an interface. RAs received on those trusted ports or linkaggs will be allowed to continue on to all other IPv6 clients reached via the interface.

Neighbor Cache Limit

The size of the neighbor cache can be limited on a system-wide basis. Once the limit is reached, no new entries will be added. The system-wide limit can be used to control the resources allocated for the IPv6 neighbor cache. A neighbor cache limit may also be specified on a per-interface basis. Once the interface's limit is reached, no new neighbor entries are allowed. The per-interface limit can be used to prevent any particular node attached to an interface from flooding the cache, either maliciously or due to a malfunction. By default, no limits are set.

Neighbor Unreachability Detection (NUD)

IPv6 Neighbor Unreachability Detection (NUD) is performed to check the status of an unconfirmed neighbor when traffic is forwarded to it. By default, up to three neighbor solicitations are sent, with an interval of one second, to reconfirm that the neighbor is reachable. In certain situations (e.g. high traffic loads), the default settings may not be sufficient to maintain the neighbor cache in a stable state. In such situations both the maximum number of neighbor solicitations and the interval at which they are sent may be modified.

ECMP
The OmniSwitch supports 16 ECMP routes for IPv6.

Border Gateway Protocol (BGP) 4

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.

IPv6 DHCP Relay

The Alcatel-Lucent OmniSwitch implementation of RFC 3315 contains DHCP relay support for IPv6.

The DHCPv6 Relay on OmniSwitch processes and forwards all DHCPv6 messages triggered by DHCPv6 client to the configured DHCPv6 relay agent as a unicast packet. A DHCPv6 relay agent is required in situations where DHCPv6 clients do not reside on the same link as the DHCP server. DHCPv6 Relay is a per-interface option that can be enabled on any IPv6 interface.

It supports multicast-capable IPv6 interfaces (VLAN and configured tunnel interfaces) and non-multicast-capable IPv6 interfaces (6to4 tunnel). The DHCPv6 Relay agent is part of the link-scoped multicast group (FF02::1:2) on the interface. Any messages sent by a client to that address will then be handled by DHCPv6 Relay agent. A maximum of five unicast or link-scoped multicast relay destinations can be configured for each interface on which DHCPv6 Relay is enabled.

Intermediate System to Intermediate System (IS-IS) for IPv6

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.

Open Shortest Path First (OSPF) v3

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. Different interface types can be configured based on the network type connected. The OmniSwitch supports point-to-point, point-to-multipoint, non-broadcast multiple access (NBMA), and broadcast interface types. 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.

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. 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 QoS software and embedded virtual output queue (VOQ) 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.

The implementation of VOQ integrates traffic management with QoS scheduling. Embedded profiles apply the QoS admission control and bandwidth management configurations to VOQ flows. Packets received by the switch are queued on the ingress to avoid congestion on the egress. A centralized scheduler in the switch fabric arbitrates flows between ingress and egress ports based on feedback from the egress port.

Auto-Qos Prioritization of 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.

Auto-QoS Prioritization of IP Phone Traffic

This feature is used to automatically enable the prioritization of IP phone traffic. The traffic can be assigned a priority value or, if set to trusted mode, the IP phone packet is used to determine the priority. IP phone traffic is identified by examining the source MAC address of the packet received on the port. If the source MAC falls within the Alcatel-Lucent ranges below, the Auto-QoS feature automatically sets the priority.

- 00-80-9F-xx-xx-xx and 00:13-FA-xx-xx-xx

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.

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.

VFC/VoQ Profiles

- Four pre-defined Qset Profiles (QSPs) are supported.
- WRED is not supported.

Multicast Feature Support

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 gleaning). 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.

Distance Vector Multicast Routing Protocol (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 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.

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.

All modes of PIM are supported for both IPv4 and IPv6. 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).

PIM/DVMRP Interoperability

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

MBR functionality is configured and enabled on OmniSwitches that are located at points where PIM and DVMRP regions interconnect. An MBR first pulls down packets generated within the PIM domain and injects them into the DVMRP domain. Then the MBR imports packets generated within the DVMRP domain so that they can be delivered to group members inside the PIM domain, using PIM mechanisms. In the case of transit networks, the MBR acts to pass the multicast traffic through both the PIM and DVMRP domains.

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

Monitoring and Troubleshooting Feature Support

Digital Diagnostic Monitoring (DDM)

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.

Fault Propagation and Link Monitoring

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

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

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

Link Fault Propagation

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

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

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

Ethernet OAM

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

Wait to Shutdown

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

An interface is administratively brought down.

An interface is shutdown from a violation.

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

Gigaword Packet Counters

Acct-Input-Octets (type-42) and Acct-Output-Octets (type-43) are sent to the RADIUS Server in accounting packets. These statistics are used by the service providers for billing of users. As these two fields are 4 bytes longer as per the RADIUS standard, it can support a maximum value of 4GB (2^32 -1= 4,294.967,295). Whenever a user uses more than 4GB, the exact count of usage is lost.

Acct-Input-Gigawords (type-52) and Acct-Output-Gigawords (type-53) attributes are introduced to overcome the limitation due to the 4 bytes size of Acct-Input-Octets and Acct-Output-Octets. These attributes indicates how many times the Acct-Input-Octets and Acct-Output-Octets counter has wrapped the 4GB traffic over the course the service being provided.

Whenever the input octets and output octets exceeds 2^32-1 bytes, before sending accounting packet to the RADIUS Server, these octets are converted into multiples of 4GB and will be sent in Acct-Input-Gigawords (type-52) and Acct-Output-Gigawords (type -53) attributes. For every 4GB traffic, the value is incremented and the remaining amount of traffic is displayed in Acct-Input-Octets and Acct-Output-Octets attribute.

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.

Line Diags & Enhanced Port Performance (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 parter, enable EPP.
- Not all transceivers support enabling EPP.

Ping and Traceroute / Extended Ping and Traceroute

Ping and Traceroute support both IPv4 and IPv6 along with additional parameters such as a source interface and timeout. The extended Ping & Traceroute functionality allows for the following additional parameters:

- Ping- Set the Source IP, Set TOS value, Set DF bit in IP header, Set data pattern, Set sweep range
- Traceroute - Set the Source IP, Set Timeout in seconds, Set Probe count, Set Min and Max TTL

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 transmited and 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 guidelines:

- 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 features 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.

Time Domain Reflectometry (TDR)

Time Domain Reflectometry (TDR) is a feature that is used to detect cable faults. This feature is best deployed in networks where service providers and system administrators want to quickly diagnose the state of a cable during outages, before proceeding with further diagnosis. When a TDR test is initiated, a signal is sent down a cable to determine the distance to a break or other discontinuity in the cable path. The length of time it takes for the signal to reach the break and return is used to estimate the distance to the discontinuity. TDR is an on-demand, out-of-service test. The test is not automatically triggered; data and protocol traffic is interrupted. Only supported on copper ports.

Uni-Directional Link Detection (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.

Metro Ethernet Feature Support

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.

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.

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

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

Application Monitoring (Appmon)

The OmniSwitch Appmon feature detects and identifies remote applications by scanning the payload of IP packets and comparing the payload to pre-defined bit patterns (application signatures). Once an application is identified, Appmon collects and stores information about the application flow in a database on the local switch. Flow identification is done based on the following 5-tuple match of the IP packets:

- Source IP Address (IPv4 or IPv6)
- Destination IP Address (IPv4 or IPv6)
- Source L4 port
- Destination L4 port

- IP Protocol (TCP or UDP)

The following are the key components for application monitoring:

- Application Pool - The pool of applications that are available for monitoring.
- Application List - The list of application names currently being monitored. Applications can be added to this list using the application name or application group.
- Application Group - A group of application names. The group can be added to the application list in place of using the application name.
- Application Record - Used to display current-hour application information as well the historical hourly or 24-hour application-records.

Note: AppMon is only supported on OmniSwitch 6860E platforms. It is supported in a virtual chassis containing OmniSwitch 6860s if at least one switch in the virtual chassis is an OmniSwitch 6860E.

The following applications are supported in this release.

- BitTorrent
- Citrix ICA HDX
- DNS
- FTP
- Jabber
- LinkedIn
- NFS
- NTP
- POP3
- SIP
- TFTP
- Twitter
- WebEx
- Windows Media Player
- Youtube

Access Control Lists (ACLs) for IPv4/IPv6

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.
- **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.

Access Guardian 2.0

Access Guardian 2.0 refers to the following security functions that work together to provide a dynamic, proactive solution to network security:

Universal Network Profile (UNP) – Network access controls are configured through the Universal Network Profile (UNP) feature to provide the following authentication and classification methods:

- MAC-based authentication and 802.1X-based authentication using a RADIUS-capable server.
- Redirection for Captive Portal authentication.
- Redirection to ClearPass Policy Manager (CPPM) for Bring Your Own Devices (BYOD) user device registration, integrity check, UNP assignment, and policy list assignment.
- Switch-wide classification rules to classify users based on port and device attributes (for example, source port, source MAC, Group ID, IP address). No authentication required.
- Default UNP classification for traffic not classified through other methods.

Basically, UNP provides a method for dynamically assigning network devices to VLAN domains. A profile consists of configurable attributes. When a device sends traffic that matches these attributes, the device is then assigned to a VLAN associated with the UNP. The UNP may also specify role-based attributes, such as a QoS/ACL policy list, a location-based policy, or a time-based policy, that are subsequently applied to device traffic associated with the UNP.

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 assignment.

Authentication, Authorization, and Access (AAA) – AAA provides the authentication and accounting functionality for 802.1x, MAC, and Captive Portal authentication by interacting with an external RADIUS-capable server.

- Up to 4 authentication servers can be configured each for 802.1x, MAC, and Captive Portal authentication. AAA tries to connect to the first AAA server configured; if the server is unreachable, access to the next AAA server is attempted.
- Up to 4 accounting servers can be configured each for 802.1x, MAC and Captive Portal accounting. Local accounting to syslog is also supported.
- The OmniSwitch acts as the authenticator for wired MAC, 802.1x, and Captive Portal-based authentication. All RADIUS messages are initiated by the OmniSwitch

Captive Portal—Internal and external Captive Portal Web-based authentication. The OmniSwitch presents default or customized Web pages to the user through an internal Web server on the switch. A post-authentication and/or post-classification process to validate user credentials and dynamically assign a new role (policy list) to enforce user access to the network. External, guest Captive Portal authentication is provided through the Access Guardian interaction with the ClearPass Policy Manager.

Quarantine Manager and Remediation (QMR) – A client MAC address is determined to be in a quarantined state when one of the following occurs:

- The OmniVista Quarantine Manager (OVQM) application receives a TRAP indicating that the MAC address has to be quarantined. The TRAP may come from a network anomaly detection application or from an IDS running in the same subnet as the client.
- A list containing the quarantined MAC address is manually configured on OVQM.

- A list containing the quarantined MAC address is manually configured on every switch in the network.

After the list of quarantined MAC addresses is known, OVQM can add these addresses to the Quarantine MAC group and push the configuration to the switches in a logical group or to all switches.

The Access Guardian Quarantine Manager and Remediation (QMR) feature moves the users associated with the quarantined MAC addresses to a QMR restricted role. A built-in policy list is associated with the QMR role that restricts quarantined users to communicating with a designated remediation server until their quarantined status is corrected.

QMR works on UNP and non-UNP ports. On non-UNP ports, L2 source learning receives the quarantined MAC addresses from QMR and changes the MAC status to 'quarantined'. However, on UNP ports, L2 source learning will not show the quarantined MAC addresses as 'quarantined'. In this case, the appropriate status of the MAC address is displayed through UNP commands.

The following QMR components are configured through Access Guardian CLI commands:

- Quarantined MAC address group. This is a reserved MAC address group that contains the MAC addresses of clients that OVQM has quarantined and that are candidates for remediation. The default name of this group is "Quarantined", but the name can be changed if necessary.
- Remediation server and exception subnet group. When a client is quarantined, all the traffic from the client is blocked by default. However, the administrator can configure access to some exception subnets to which the quarantined client can be redirected, such as the IP address of a remediation server to obtain updates and correct its quarantined state. Configuring a maximum of three subnets is allowed.
- A remediation server URL to which quarantined clients are redirected.
- Quarantined Page. When a client is quarantined and if a remediation server URL is not configured, QMR can send a Quarantine Page to notify the client of its quarantined state.
- QMR custom proxy port. This specifies an HTTP proxy port number to which quarantined client traffic is redirected for remediation. The default HTTP ports used are TCP 80 and TCP 8080.

Bring Your Own Device (BYOD) - OmniSwitch / ClearPass (CPPM 6.3) Integration: Guest users and user devices information can be allowed to access specific network resources. BYOD support provides restricted access to the network so that the end user device can be validated, user roles identified, compliance checked, and have the correct access policies applied. The OmniSwitch leverages the Access Guardian features along with the ClearPass Policy Manager to provide the overall BYOD solution. This feature supports the following functionalities:

- Unified access policy management solution for Wireline and Wireless networks using CPPM
- Integration with Access Guardian UNPs and MAC/802.1x authentication
- Restricts access to the network and validation for end user devices including employees with IT supplied devices, IP phones, employees personal devices, guest devices, access points, cameras, and silent devices such as printers.
- CPPM can act as a RADIUS server for new deployments or RADIUS proxy for existing networks. Self-service/self-registration by Employees when they connect to the Enterprise network using their personal device through CPPM.
- Captive portal hosted on CPPM for this feature.
- Device Profiling and Posture Check. Registration and tracking of devices associated with Employees and approved for usage.
- Redirection and restricted access for non-compliant devices.
- Zero-touch Auto-configuration of employee personal devices based on pre-defined role-based Configuration profiles.

- Differentiated access & user experience policies based on Corporate or Employee Personal device, Applications and Role.
- Integration with RADIUS Server and CPPM for Authentication, Authorization and Accounting.
- Automatic provisioning of Applications such as NAC Agent, MDM Client as part of the device enrollment process on Employee Personal Devices.
- Automatic provisioning of Device Certificates that are dynamically requested, issued and installed on the Employee Personal Device with association to Employee corporate Credentials
- Provides notification of BYOD policy violations, usage statistics, time and cost information to the end-user in real-time.
- RADIUS Change of Authorization (CoA)
    - A mechanism to change AAA attributes of a session after authentication
    - New Profile sent as an attribute in the message
    - Disconnect Message to terminate user session and discard all user context
    - Port bounce capability can be configured on the OmniSwitch to ensure a clean re-authentication process for non-supplicant devices.

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 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 leaned.

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.

Early Availability / Demonstration Feature Descriptions

The following software features are being introduced with the 8.1.1.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.

Deep Packet Inspection (DPI)

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

- Monitoring – No action, accounting and visibility
- QoS / UNP – If there is a match, applies QoS or UNP policy on matching flow

The following applications are supported for demonstration purposes in this release.

- Jabber, YouTube, WebEx, Citrix ICA HDX, NFS, NTP, SIP, POP3, Windows Media Player

SNMP Traps

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

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

No.	Trap Name	Platforms	Description
16	bgpEstablished	All	The BGP routing protocol has entered the established state.
17	bgpBackwardTransition	All	This trap is generated when the BGP router port has moved from a more active to a less active state.
18	esmDrvTrapDropsLink	All	This trap is sent when the Ethernet code drops the link because of excessive errors.
19	portViolationTrap	All	This trap is sent when a port violation occurs. The port violation trap will indicate the source of the violation and the reason for the violation.
20	dvmrpNeighborLoss	All	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	All	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	All	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	All	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	All	Sent by a bridge that became the new root of the spanning tree.

No.	Trap Name	Platforms	Description
25	stpRootPortChange	All	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	All	This trap is sent when any NI fails to configure mirroring. Due to this error, port mirroring session will be terminated.
27	mirrorUnlikeNi	All	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	All	A change occurred in the operational status of the server load balancing entity.
29	sessionAuthenticationTrap	All	An authentication failure trap is sent each time a user authentication is refused.
30	trapAbsorptionTrap	All	The absorption trap is sent when a trap has been absorbed at least once.
31	alaDoSTrap	All	Indicates that the sending agent has received a Denial of Service (DoS) attack.
32	ospfNbrStateChange	All	Indicates a state change of the neighbor relationship.
33	ospfVirtNbrStateChange	All	Indicates a state change of the virtual neighbor relationship.
34	lnkaggAggUp	All	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	All	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	All	This trap is sent when any given port of the link aggregate group goes to the attached state.
37	lnkaggPortLeave	All	This trap is sent when any given port detaches from the link aggregate group.
38	lnkaggPortRemove	All	This trap is sent when any given port of the link aggregate group is removed due to an invalid configura tion.
39	monitorFileWritten	All	This trap is sent when the amount of data requested has been written by the port

No.	Trap Name	Platforms	Description
			monitoring instance.
40	alaVrrp3TrapProtoError	All	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
41	alaVrrp3TrapNewMaster	All	The SNMP agent has transferred from the backup state to the master state.
42	chassisTrapsPossibleDuplicateMac	All	The old PRIMARY element cannot be detected in the stack. There is a possiblity of a duplicate MAC address in the network
43	lldpRemTablesChange	All	A lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes.
44	pimNeighborLoss	All	A pimNeighborLoss notification signifies the loss of an adjacency with a neighbor.
45	pimInvalidRegister	All	An pimInvalidRegister notification signifies that an invalid PIM Register message was received by this device
46	pimInvalidJoinPrune	All	A pimInvalidJoinPrune notification signifies that an invalid PIM Join/Prune message was received by this device.
47	pimRPMappingChange	All	An pimRPMappingChange notification signifies a change to the active RP mapping on this device.
48	pimInterfaceElection	All	An pimInterfaceElection notification signifies that a new DR or DR has been elected on a network.
49	pimBsrElectedBSRLostElection	All	This trap is sent when the current E-BSR loses an election to a new Candidate-BSR.
50	pimBsrCandidateBSRWinElection	All	This trap is sent when a C-BSR wins a BSR Election.
51	lpsViolationTrap	All	A Learned Port Security (LPS) violation has occurred.
52	lpsPortUpAfterLearningWindowExpiredT	All	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	All	Generated when an LPS port learns a bridged MAC.
54	gvrpVlanLimitReachedEvent	All	Generated when the number of vlans learned dynamically by GVRP has reached

No.	Trap Name	Platforms	Description
			a configured limit.
55	alaNetSecPortTrapAnomaly	All	Trap for an anomaly detected on a port.
56	alaNetSecPortTrapQuarantine	All	Trap for an anomalous port quarantine.
57	ifMauJabberTrap	All	This trap is sent whenever a managed interface MAU enters the jabber state.
58	udldStateChange	All	Generated when the state of the UDLD protocol changes.
59	ndpMaxLimitReached	All	This IPv6 Trap is sent when the hardware table has reached the maximum number of entries supported.
60	ripRouteMaxLimitReached	All	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	All	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	All	This trap is sent when the ERP Ring State has changed from “Idle” to “Protection”.
63	alaErpRingMultipleRpl	All	This trap is sent when multiple RPLs are detected in the Ring.
64	alaErpRingRemoved	All	This trap is sent when the Ring is removed dynamically.
65	ntpMaxAssociation	All	This trap is generated when the maximum number of peer and client associations configured for the switch is exceeded.
66	ddmTemperatureThresholdViolated	All	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	All	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	All	This trap is sent when if an SFP/ XFP/SFP+ Tx bias current has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex.

No.	Trap Name	Platforms	Description
			It also provides the current realtime value of SFP/XFP/SFP+ Tx bias current.
69	ddmTxPowerThresholdViolated	All	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	All	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	All	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	Not Supported	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is up.
73	multiChassisIpcVlanDown	Not Supported	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is down.
74	multiChassisMisconfigurationFailure	Not Supported	This trap is sent to indicate a mis-configuration due to Chassis Id or IPC VLAN.
75	multiChassisHelloIntervalConsisFailure	Not Supported	This trap is sent to indicate a hello interval consistency failure.
76	multiChassisStpModeConsisFailure	Not Supported	This trap is sent to indicate an STP mode consistency failure.
77	multiChassisStpPathCostModeConsisFailure	Not Supported	This trap is sent to indicate an STP path cost mode consistency failure.
78	multiChassisVflinkStatusConsisFailure	Not Supported	This trap is sent to indicate a VFLink status consistency failure.
79	multiChassisStpBlockingStatus	Not Supported	This trap is sent to indicate the STP status for some VLANs on the VFLink is in blocking state.
80	multiChassisLoopDetected	Not Supported	This trap is sent to indicate a loop has been detected.
81	multiChassisHelloTimeout	Not Supported	This trap is sent to indicate the hello timeout has occurred.
82	multiChassisVflinkDown	Not Supported	This trap is sent to indicate the VFLink is

No.	Trap Name	Platforms	Description
			down.
83	multiChassisVFLMemberJoinFailure	Not Supported	This trap is sent to indicate a port configured as virtual-fabric member is unable to join the virtual-fabric link.
84	alaDHLVlanMoveTrap	All	When linkA or linkB goes down or comes up and both ports are are part of some vlan-map, this trap is sent to the Management Entity, with the DHL port information.
85	alaDhcpClientAddressAddTrap	All	This trap is sent when a new IP address is assigned to DHCP Cli-ent interface.
86	alaDhcpClientAddressExpiryTrap	All	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	All	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	All	This notification is generated when the system enters or leaves the overload state.
89	vRtrIsisManualAddressDrops	All	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
90	vRtrIsisCorruptedLSPDetected	All	This notification is generated when an LSP that was stored in memory has become corrupted.
91	vRtrIsisMaxSeqExceedAttempt	All	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
92	vRtrIsisIDLenMismatch	All	A notification sent when a PDU is received with a different value of the System ID Length.
93	vRtrIsisMaxAreaAddrsMismatch	All	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
94	vRtrIsisOwnLSPPurge	All	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
95	vRtrIsisSequenceNumberSkip	All	When an LSP is received without a System ID and different contents.
96	vRtrIsisAutTypeFail	All	A notification sent when a PDU is received with the wrong authentication type field.
97	vRtrIsisAuthFail	All	A notification sent when a PDU is received with an incorrent authentication information field.

No.	Trap Name	Platforms	Description
98	vRtrIsisVersionSkew	All	A notification sent when a a Hello PDU is received from an IS running a different version of the protocol.
99	vRtrIsisAreaMismatch	All	A notification sent when a Hello PDU is received from an IS which does not share any area address.
100	vRtrIsisRejectedAdjacency	All	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	All	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
102	vRtrIsisOrigLSPBufSizeMismatch	All	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	All	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
104	vRtrIsisAdjacencyChange	All	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	All	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	All	A notification sent when an adjancency's graceful restart status changes.
107	mvrpVlanLimitReachedEvent	All	This trap is sent when the num-ber of VLANs learned dynamically by MVRP reaches the configured limit.
108	alaHAVlanClusterPeerMismatch	Not Supported	The trap is sent when parameteras configured for this cluster ID (Level 1 check) does not match accross the MCLAG peers.

No.	Trap Name	Platforms	Description
109	alaHAVlanMCPeerMismatch	Not Supported	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	All	The trap is sent when the dynamic MAC is learned on non-server cluster port
111	unpMcLagMacIgnored	Not Supported	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	Not Supported	This trap is sent when a configuration becomes “Out of Sync".
113	multiChassisGroupConsisFailure	Not Supported	This trap is sent when there is an inconsistency between local and peer chassis group.
114	multiChassisTypeConsisFailure	Not Supported	This trap is sent when there is an inconsistency between local and peer chassis group.
115	alaPimNonBidirHello	All	This trap is sent when a bidir-capable router has received a PIM hello from a non-bidir-capable router. It is generated whenever the counter alaPimsmNon-BidirHelloMsgsRcvd is incremented, subject to the rate limit specified by alaPimsmNonBidirHelloNotificationPeriod.
116	dot1agCfmFaultAlarm	All	This trap is sent when a MEP has a persistent defect condition. A notification (fault alarm) is sent to the management entity with the OID of the MEP that has detected the fault.
117	alaSaaIPIterationCompleteTrap	All	This trap is sent when an IP SAA iteration is completed.
118	alaSaaEthIterationCompleteTrap	All	This trap is sent when when an eth-LB or Eth-DMM SAA iteration is completed.
119	alaSaaMacIterationCompleteTrap	All	This trap is sent when a MAC iteration is complete.
120	virtualChassisStatusChange	All	This trap is sent when a chassis status change is detected.
121	virtualChassisRoleChange	All	This trap is sent when a chassis role change is detected.
122	virtualChassisVflStatusChange	All	This trap is sent when s vflink status change is detected.

No.	Trap Name	Platforms	Description
123	virtualChassisVflMemberPortStatusCh	All	This trap is sent when a vflink member port has a change of status.
124	virtualChassisVflMemberPortJoinFail	All	This trap is sent when a port configured as virtual-fabric member is unable to join the virtual-fabric link.
125	lldpV2RemTablesChange	All	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	All	This trap is sent when the LDP module changes state either administratively or operationally.
127	evbFailedCdcpTlvTrap	All	This trap is sent when bridge receives a CDCP packet with: - Wrong TLV type, or - Wrong OUI, or - Role is set to Bridge, or - Wrong default channel(scid), or - Incorrect channel number(scid).
128	evbFailedEvbTlvTrap	All	This trap is sent when bridge receives an EVBTLV packet with: - Wrong TLV type. or - Incorrect TLV length, or - Wrong OUI.
129	evbUnknownVsiManagerTrap	All	This trap is sent when bridge receives a VDP packet with: - Unknown Manager ID type, or - Wrong Manager ID length.
130	evbVdpAssocTlvTrap	All	This trap is sent when bridge receives an ASSOC TLV in a VDP packet with: - Null VID found and number of entry field is not 1, or - Unknown filter format, - Null VID on De-Assoc TLV type, or - VSI included more than Max number of filter info entries

No.	Trap Name	Platforms	Description
131	evbCdcpLIdpExpiredTrap	All	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	All	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	All	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	All	This trap is sent when there is a failure to create/delete a service.
135	smgrServiceHwError	All	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	All	This trap is sent when there is a failure to create/delete a Service Access Point.
137	smgrServiceSapHwError	All	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	All	This trap is sent when there is a failure to create/delete a Service Distribution Point.
139	smgrServiceSdpHwError	All	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	All	This trap is sent when there is a failure to create/delete an SDP Bind.
141	smgrServiceSdpBindHwError	All	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	All	This trap is sent when there is a .general system failure detected during normal system operation.
143	smgrStatusChange	All	This trap is sent when there is a status change for a group of selected services.
144	portViolationNotificationTrap	All	This trap is sent when a port violation is cleared.
145	multiChassisConsisFailureRecovered	Not Supported	This trap is sent when the system has recovered from a multi-chassis inconsistency between the local and peer switches

No.	Trap Name	Platforms	Description
146	alaSaaPacketLossTrap	All	This trap is sent when a a packet is lost during a test.
147	alaSaaJitterThresholdYellowTrap	All	This trap is sent when the Jitter Threshold crosses 90%.
148	alaSaaRTTThresholdYellowTrap	All	This trap is sent when the RTT Threshold crosses 90%.
149	alaSaaJitterThresholdRedTrap	All	This trap is sent when the Jitter threshold is crossed.
150	alaSaaRTTThresholdRedTrap	All	This trap is sent when the RTT threshold is crossed.
151	chassisTrapsDuplicateMacClear	All	This trap is sent when the old Master Chassis has rejoined the Virtual Chassis as a slave.
152	alaFipResourceThresholdReached	All	The allowed maximum percentage of filter resources configured from the allocated FIPS resources is exceeded.
153	virtualChassisUpgradeComplete	All	Critical trap indicates whether the software upgrade process has failed after a timeout or completed successfully. Note that if the process fails, it may be still possible for the system to recover if the process successfully completes later after the expired timeout.
154	appFPSignatureMatchTrap	All	This trap is sent when a traffic flow matches an application signature.
155	virtualChassisVflSpeedTypeChange	All	This trap is sent when the VFL speed type is changed.
156	alaSIPSnoopingACLPreemptedBySOSCall	All	This trap is sent when a SIP snooping RTP/RTCP ACL entry is preempted by an SOS call.
157	alaSIPSnoopingRTCPOverThreshold	All	This trap is sent when one or more RTCP parameters exceeds the threshold limit.
158	alaSIPSnoopingRTCPPktsLost	All	This trap is sent when RTCP packets are lost due to rate limiting.
159	alaSIPSnoopingSignallingLost	All	This trap is sent when SIP signaling messages are lost due to rate limiting.
160	alaSIPSnoopingCallRecordsFileMoved	All	This trap is sent when the SIP Snooping Ended Call Records flash file is moved from /flash/switch/sip_call_record.txt to /flash/switch/sip_call_record.txt.old. This happens when the configured call record storage limit is reached and possibly at boot-up if /flash/switch/sip_call_record.txt from previous run exists at the first check.

No.	Trap Name	Platforms	Description
161	alaIPv6NeighborLimitExceeded	All	
162	alaIPv6NeighborVRFLimitExceeded	All	
163	alaIPv6InterfaceNeighborLimitExceed	All	
164	alaDyingGaspTrap	All	This trap is sent when a switch has lost all power.
165	alaDhcpSrvLeaseUtilizationThreshold	All	This trap is sent when the lease utilization on a subnet exceeds or falls below the configured threshold value.
166	alaDHCPv6SrvLeaseUtilizationThresho	All	This trap is sent when the lease utilization on a subnet exceeds or falls below the configured threshold value.
167	smgrServiceStatusChange	All	This trap is sent when there is a change in service operating status. A service is operationally up when it's admin-up and there's at least one active SAP or one active bind that is operationally up.
168	smgrSapStatusChange	All	This trap is sent when there is a change in SAP operating status. A SAP is operationally up when it's admin-up and the link status of the physical or logical port of the SAP is operationally up.
169	smgrSdpStatusChange	All	This trap is sent when there is a change in SDP operating status. For SPB, the SDP is dynamically created or destroyed as calculated by ISIS protocol when a unicast/multicast path to reach a neighbor node is determined.
170	smgrSdpBindStatusChange	All	This trap is sent when there is a change in SDP Bind operating status. For SPB, the SDP Bind is dynamically created or destroyed as detected by ISIS when the same ISID is configured in the neighbor node.
171	alaPethPwrSupplyConflictTrap	All	This trap is sent when there is a power supply type conflict.
172	alaPethPwrSupplyNotSupportedTrap	All	This trap is sent when the power supply is not supported.
173	chasTrapsBPSLessAllocSysPwr	All	This trap is sent when there is insufficient system power provided by the BPS.
174	chasTrapsBPSStateChange	All	This trap is sent when BPS Power Supplies are inserted or removed.
175	chasTrapsNiBPSFETStateChange	All	This trap is sent when there is a BPS FET

No.	Trap Name	Platforms	Description
			change of state.
176	alaDhcpBindingDuplicateEntry	All	This trap is sent when there is a MAC movement in the DHCP-Binding Table.
177	alaVCSPProtectionTrap	All	This trap is sent when a virtual chassis enters into VCSP Protection state
178	alaVCSPRecoveryTrap	All	This trap is sent when a virtual chassis enters into VCSP Active state
179	pethPsePortOnOffNotification	All	This trap is sent to indicate whether or not the PSE Port is delivering power to the PD. This notification SHOULD be sent on every status change except in the searching mode.
180	pethMainPowerUsageOnNotification	All	This trap is sent to indicate that the PSE Threshold usage indication is on, the usage power is above the threshold. At least 500 msec must elapse between notifications being emitted by the same object instance.
181	pethMainPowerUsageOffNotification	All	This trap is sent to indicate that the PSE Threshold usage indication is off, the usage power is below the threshold. At least 500 msec must elapse between notifications being emitted by the same object instance.
182	chasTrapsBPSFwUpgradeAlert	All	This trap is sent when a BPS firmware upgrade is required.
183	alaAppMonAppRecordFileCreated	All	This trap is sent when the application records monitored in the past hour are written to the flash file.
184	alaAppMonFlowRecordFileCreated	All	This trap is sent when a pre-configured number of flow records (configured by setting alaAppMonLoggingThresholdFlows) is written to the flash file.
185	alaDPIFlowRecordFileCreated	All	This trap is sent when a pre-configured number of flow records (configured by setting alaDPILoggingThresholdFlows) is written to the flash file.

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
SAA	All	N/A
Internal DHCP Server	All	N/A
Traffic Anomaly Detection (Network Security)	All	N/A
Multi-Chassis Link Aggregation (MC-LAG)	All	N/A
IPv4 over SPB	All	N/A
DPI (EA/Demo Feature)	All	N/A

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
SAA	All 'saa' commands
Network Security	All 'netsec' commands
Internal DHCP Server	dhcp-server {restart | enable | disable} dhcpv6-server {restart | enable | disable}
MC-LAG	All 'multi-chassis' commands
IPv4 over SPB	'spb ipvpn bind vrf'
DPI (EA/Demo Feature)	All DPI related CLI and parameters

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.

Access Guardian / UNP

PR	Description	Workaround
192012	RADIUS test tool cannot send calling-station-ID in test request for CPPM to identify and return redirect URL.	There is no known workaround at this time.
192486	Redirection may not work with Chrome browser when opened with https://.	Enter a URL that does not use HTTPS.
194300	An IP logical interface bound to a physical port will be disabled after enabling UNP on the same port.	This configuration is not allowed. Remove UNP on the port to bring the IP logical interface back up.

IPMS

PR	Description	Workaround
194456	The ipms static-neighbor and static-querier feature does not capture a membership report that is sent by another client (interested in the same group) in response to the group specific query sent by the querier which was sent in response to a leave group message by the other client.	There is no known workaround at this time.

PoE

PR	Description	Workaround
192825	Port range is not supported in the following CLI command: "lanpower port {admin-state | priority | power | type}	Configure ports individually.
192875	When two power supplies are connected (920W/600W) and power draw is more than 450W for 600W power supply and more than 780W for a 920W power supply, lanpower for one of the ports (lowest priority port) resets when the AC power cord is removed and re-inserted on any one of the power supplies.	If the power consumption is within the supported 450W for a 600W power supply and within 780W for a 920W power supply, this issue will not be seen.
194461	In a system with two 920W power supplies, when one power supply AC cord is removed and re-inserted, the DC LED	There is no functional impact. However, to clear the issue perform a power supply extraction and insertion with the power

	for the other power supply may begin blinking.	supply which does have the DC LED blinking. **Note**: AC cord should always be removed prior to extracting power supplies.

Layer 3

PR	Description	Workaround
189973	When an IS-IS VLAN is configured with multiple subnets, adjacency will not form after a toggle of the primary ip interface.	Use the command 'ip interface <interface name> primary' to ensure ISIS adjacency is formed when a primary IP interface is toggled.
188749	Multiple DHCP release messages are sent by the relay agent.	There is no known workaround. There is no functional impact.
192203	dhcp-snooping hardware resource limits are reached in a VC which has IP source filtering enabled on a linkagg, when one of the VC units is reloaded.	There is no known workaround at this time.
193501	The 'ip udp relay port <num> description <description>' command does not save the description.	There is no known workaround at this time.
190258	There may be a discrepency between the number of VLAN or port ingress source filtering binding entries learned vs. the number supported.	The following ISF entries are supported based on the configuration: 32 VLAN ISF enabled - support for 192 ISF entries Port ISF + 32 VLAN ISF enabled - support 190 ISF entries With only port ISF enabled - support 254 entries

QoS

PR	Description	Workaround
194133	All ports are untrusted by default, and adding a tagged VLAN member to a port does not automatically make the ports trusted.	Manually configure the port as trusted.

Security

PR	Description	Workaround
191476	YouTube video traffic will not be recognized by application monitoring if it is encrypted.	There is no known workaround at this time.

WebView

PR	Description	Workaround
186561	Firefox 23 and previous versions can't access WebView over an IPv6 interface.	Upgrade to version 24 or higher or use Internet Explorer.
189471	Unable to delete LPS filtered macs from webview.	Use cli command 'mac-learning flush dynamic mac-address <mac-address>'
191880	When more than 100 MAC are learned on a UNP port they may not be displayed in WebView.	Use the CLI to search for the specific mac: -> show mac-learning mac-address <mac-address>

Virtual Chassis

PR	Description	Workaround
192519	Not all show commands executed on a slave chassis that is in a shutdown state will show an accurate representation of the actual state or status.	There is no known workaround at this time.
194180	VFL ports display as untrusted in "show qos ports" output.	VFL ports are always trusted. This is a display issue only.
193578	High CPU utilitzation may be seen on slave chassis after multiple takeovers.	There is no known workaround at this time.
194237	Image validation failure seen on slave chassis during RCL process when vcsetup.cfg file is present.	Remove the vcsetup.cfg file before performing the RCL process.

System / Monitoring

PR	Description	Workaround
193380	Console output may be garbled after a power-cycle.	Disconnect the cable connected to the console port (RS232). Power up the unit and wait for 2 seconds. Re-connect the cable to the console port.
191017	TDR statistics are not maintained after a VC takeover.	Obtain the statistics again after the VC takeover.
191935	When port-mirroring is enabled on a 10G NNI port up to 25% of packets may not be mirrored.	Disable and re-enable port-mirroring session or toggle the destination mirroring port.
194447	SLB server down notification messages are sometimes seen after the completion of a vc-takeover.	There is no known workaround at this time. This is a display issue only.

Hot Swap Guidelines

Hot Swap Feature Guidelines

- Hot swap of like power supplies is supported.
- Hot swap of unlike power supplies is not supported.
- Hot insertion, the insertion of a power supply into a previously empty slot, is supported.
- Mixing of different wattage power supplies in the same chassis is not supported.

Hot Swap Procedure

The following steps must be followed when hot-swapping power supplies.

1. Disconnect the power supply cord from the power supply.
2. Extract the power supply.
3. Insert replacement power supply of same type.
4. Connect the power supply cord to the new power supply.

Technical Support

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

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

Email: esd.support@alcatel-lucent.com

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

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

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

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

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

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

Third Party Licenses and Notices

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

Appendix A – 6.X to 8.X Feature Comparison Summary

The table below is a a feature comparison between 6.X and 8.X. In most cases the 6.X release in which a feature was introduced has also been indicated.

Feature	Introduced	8.X
Hardware/Stacking Features:		
Virtual Chassis (Stacking) of 8 switches	6.X	Supported (VC)
DDM - Transceiver Digital Diagnostic Monitoring	6.4.2	Supported
Remote Stacking	6.4.2	Supported
USB Support	6.4.3	Supported
Auto Negotiation of PoE Class	6.4.4	Supported
802.3at PoE support	6.4.4	Supported
OmniSwitch Backup Power Shelf (BPS)	6.4.5	Supported
IEEE 802.3ah Dying Gasp	6.4.5	Supported
ISSU in Stacking Configuration	6.4.5	Supported
Split Stack Protection (SSP)	6.4.6	Supported
LLDP PoE power negotiation	6.4.6	Not Supported
802.3az Energy Efficient Ethernet (EEE)	6.X	Supported
Management Features :		
Text File Configuration	6.X	Supported
TFTP Client for IPv4	6.X	Supported
Secure Copy (SCP)	6.X	Supported
Secure Shell (SSH)	6.X	Supported
Partitioned Switch Management	6.X	Supported
End User Partitioning (User profiles)	6.X	Not Supported
Extended Ping & Traceroute with choosable source IP address	6.4.3	Supported
IP Loopback0 Address In the Same Range of Existing Subnet	6.4.3	Supported
LLDP Rogue Detection	6.4.4	Not Supported
LLDP Voice Vlan Support	6.4.3	Supported
Auto-Configuration with Dynamic Management	6.4.4	Supported

Feature	Introduced	8.X
VLAN		
Auto-Confguration with Tagged Management VLAN support	6.4.5	Supported
Additional SWLOG message when link up/down event SNMP trap is sent	6.4.5	Not Supported
Ping and traceroute for read only users	6.4.5	Supported
Option to build a default user profile for admins	6.4.5	Not Supported
Per command authorization for TACACS	6.4.5	Not Supported
ssh for read only users	6.4.5	Supported
Increase system name to 254	6.4.5	Supported
Improved Captive Portal Performance	6.4.5	Supported (AG 2.0)
Enabling or Disabling Console Session	6.4.6	Not Supported
Interface Admin Down Warning	6.X	Not Supported
Layer 2 Features:		
Automatic VLAN Containment (AVC)	6.X	Supported
Multiple VLAN Registration Protocol (MVRP)	6.4.3	Supported
RRSTP	6.X	Not Supported
PVST+	6.X	Supported
Link Aggregation (static & 802.3ad)	6.X	Supported
GVRP	6.X	Not Supported
Ethernet Ring Protection (G.8032)	6.X	Supported
Ethernet Ring Protection (ERP) - Shared VLAN	6.X	Not supported
802.1Q	6.X	Supported
STP Loop Guard	6.4.4	Not Supported
Dual-Home Link (DHL)	6.4.4	Not Supported
Multi-Chassis Link Aggregation (MC-LAG)	6.4.5	Not Supported
Non-unicast Load Balancing on Link Aggregation	6.4.3	Supported
Disable MAC learning per VLAN	6.4.2	Supported
Disable MAC learning per port	6.4.2	Supported

Feature	Introduced	8.X
Layer 3 Features:		
Server Load Balancing	6.X	Supported
Recursive Static Route	6.4.3	Supported
IS-IS v4/v6	6.X	Supported
IP Multinetting	6.X	Supported
IP Route Map Redistribution	6.X	Supported
IP-IP Tunneling	6.X	Supported
Generic Routing Encapsulation (GRE)	6.X	Supported
Bi-Directional Forwarding Detection (BFD)	6.X	Supported
BGP Graceful Restart	6.X	Supported
BGP4	6.X	Supported
IGMP Relay - Forward to Specific Host in L3 Environment	6.4.2	Not supported
ECMP – Support for up to 16 paths	6.4.2	Supported
IPv6 DHCP Relay	6.4.5	Supported
Session Initiated Protocol (SIP) Snooping	6.4.5	Supported
IP interface name up to 32 character	6.4.5	Supported
Bind physical port with IP interface (routing to an IP interface)	6.4.5	Supported
Convert local interfaces into OSPF passive interfaces using route map	6.4.5	Supported
UDP port relay to specific ip-address	6.4.5	Not Supported
Automatic OSPF P2P static neighbors	6.4.5	Not Supported
M-ISIS	6.4.6	Not Supported
Monitoring/Troubleshooting Features :		
Service Assurance Agent (SAA)	6.X	Not Supported
Switch Logging	6.X	Supported
sFlow	6.X	Supported
RMON	6.X	Supported
Port Monitoring	6.X	Supported
Port Mirroring (128:1)	6.X	Supported

Feature	Introduced	8.X
Remote Port Mirroring	6.X	Supported
Link Monitoring/Flapping Detection	6.4.4	Supported
Link Fault Propogation	6.4.4	Supported
Interface Violation Recovery	6.4.4	Not Supported
Time Domain Reflectometry	6.4.4	Supported
Additional Storm Control option on AOS	6.4.5	Not Supported
Loopback Detection	6.4.5	Not Supported
Gigaword Packet Counter	6.4.6	Supported
Multicast Features:		
IPv4 Multicast Switching (IPMS)	6.X	Supported
IPv4 Multicast Switching (Proxying)	6.X	Supported
Group Address and Mask	6.X	Not supported
IPMS Flood Unknown Option	6.X	Supported
IP MC VLAN – Support for multiple sender ports	6.X	Not Supported
IGMP Multicast Group Configuration Limit	6.X	Supported
DVMRP	6.X	Supported (no DVMRP tunnels)
PIM-BFD Multicast Subsecond Convergence	6.4.5	Not Supported
Layer 2 Multicast VLAN Replication	6.4.5	Not Supported
IGMP v1/v2 to PIM-SSM Static Mapping	6.4.5	Not Supported
PIM Startup Delay	6.4.6	Not Supported
Initial Multicast Packet Routing	6.4.6	Not Supported
Multicast Address Boundaries out the 239.0.0.0/8 range	6.4.6	Not Supported
L2 Star-G mode	6.4.6	Not Supported
QoS Features :		
Port-based Ingress Limiting	6.X	Supported
Policy Based Mirroring	6.X	Supported
Policy Based Routing (Permanent Mode)	6.X	Supported
Auto-Qos Prioritization of IP Phone Traffic	6.X	Supported

Feature	Introduced	8.X
Auto-Qos Prioritization of NMS Traffic	6.X	Supported
Egress Policy Rules	6.4.3	Supported
sr-TCM and tr-TCM (RFC 2697/2698)	6.4.3	Supported
IEEE 802.1q/ad CFI/DEI Bit Stamping	6.4.3	Supported
Policy Condition Enhancements (VLAN group, 802.1p Range)	6.4.3	Supported
Flexible Inner DSCP/ToS Mapping to Outer 802.1p	6.4.3	Supported
Ingress/Egress Bandwidth via RADIUS	6.4.5	Not Supported
Per-port Rate Limiting	6.4.6	Not Supported
Security Features :		
Port Mapping	6.X	Supported
SSH Public Key Authentication	6.X	Not Supported
Learned Port Security (LPS)	6.X	Supported
BPDU Shutdown Ports	6.X	Not supported
BPDU Shutdown Auto-Recovery Timer	6.4.3	Not supported
ACL Manager (ACLMAN)	6.X	Not Supported
Account & Password Policies	6.X	Supported
ARP Spoofing Defense	6.X	Supported
ARP Poisoning Detect	6.X	Not supported
Authenticated Switch Access	6.X	Supported
Authenticated VLANs	6.X	Not Supported (Use AG 2.0)
Access Control Lists (ACLs)	6.X	Supported
802.1x Radius-down Fail-Open	6.4.2	Not supported
Admin User Remote Access Restriction Control	6.4.3	Supported
Service Type information in RADIUS Access Request	6.4.3	Supported
Allow policy list definition for HIC	6.4.5	Not Supported
SNMPv3 FIPS 140-2 Encryption Modules	6.4.5	Supported
RADIUS Test Tool	6.4.5	Supported
User Detection and domain-based	6.4.5	Not Supported

Feature	Introduced	8.X
profiles/kerberos snooping		
Virtual Network Profile	6.4.5	Not Supported
Add additional information to "show 802.1x users" command	6.4.5	Not Supported
ARP poisoning protection AOS command	6.4.5	Not Supported
Radius Calling station ID	6.4.5	Supported
802.1X on IPMVLAN	6.4.5	Not Supported
LPS sticky mode without learning windows	6.4.5	Not Supported
Case Sensitive MAC address	6.4.6	Not Supported
HIC HTTPS Web Redirection	6.4.6	Not Supported
VRF Features :		
BFD Support	6.4.2	Supported
VRRP Support	6.4.2	Supported
Switch Authentication (ASA)	6.4.2	Supported
Switch Access and Utilities	6.4.2	Supported
UDP/DHCP Relay	6.4.2	Supported
VRF Aware Multicast Routing (PIM)	6.4.3	Supported
VRF Route Leak	6.4.5	Supported (IPv4 only)
PIM SSM static map	6.4.5	Not Supported
IP Helper per-VLAN / per-VRF	6.4.6	Per VRF supported but not per VLAN. Supported on default VRF only.
Access Guardian / BYOD		
Clearpass and Access Guardian Integration Change of Authorization (CoA) Port Bounce and URL redirect	6.4.6	Supported (AG 2.0)
VDI	6.4.6	Not Supported
mDNS Relay	6.4.6	Supported
Quarantine Manager and Remediation	6.X	Supported (AG 2.0)

Feature	Introduced	8.X
Mac Authentication for Supplicant/Non-Supplicant	6.X	Supported (AG 2.0)
Dynamic VLAN Assignment (Mobility)	6.X	Not supported (Use AG 2.0)
802.1x Device Classification	6.X	Supported
802.1x RADIUS Failover	6.X	Supported
Captive Portal	6.X	Supported
Host Integrity Check (HIC)	6.X	Not Supported
Accounting for Non-supplicants	6.4.4	Supported (AG 2.0)
User Network Profiles	6.4.4	Supported (AG 2.0)
Javaless Captive Portal and MAC OS Support	6.4.3	Not supported (Use AG 2.0)
DHCP / UDP		
Per-VLAN DHCP Relay	6.X	Supported (not per VRF)
DHCP Option-82 CVLAN	6.4.4	Not Supported
DHCP Option 82 ASCII support	6.4.3	Supported
DHCP Snooping Option 82 – Port-based format	6.4.2	Supported
Internal DHCP Server	6.4.3	Not Supported
DHCP Client with configurable option 60	6.4.3	Supported
Metro Ethernet Services		
802.1q Capability on NNI ports	6.4.6	Not Supported
Virtual MEP – UNI Loopback	6.4.4	Not Supported
L2 control protocol (SW version) enhancement	6.4.5	Not Supported
CPE Test Head	6.4.5	Not Supported
Control Frame Tunnelling	6.4.5	Not Supported
PPPoE-IA	6.4.5	Not Supported
Egress Rate Limiting	6.4.2	Not supported
Ethernet OAM 802.3ah – EFM	6.4.2	Not supported
IEEE 802.1ag Version 8.1	6.4.3	Supported
ITU Y.1731	6.4.3	Supported

Feature	Introduced	8.X
Service Assurance Agent (SAA) for OAM and IP SLA Measurements	6.4.3	Not Supported
L2 Control Protocol Tunneling (L2CP)	6.4.3	Not Supported
SVLAN Routing	6.4.3	Not Supported
IPMVLAN Group Address and Mask	6.4.2	Not supported
MPLS/VPLS	6.4.2	Not Supported
Server Load Balancing – Weight Round Robin	6.4.2	Supported
Hashing Control	6.4.2	Supported
IPv6		
Unique Local IPv6 Unicast Addresses	6.X	Supported
IPv6 Scoped Multicast Addresses	6.X	Supported
IPv6 Multicast Routing	6.X	Supported
IPv6 Multicast Switching (MLD)	6.X	Supported
IPv6 Multicast Switching (Proxying)	6.X	Supported
IPv6 Client and/or Server Support	6.X	Supported
IPv6 Routing	6.X	Supported
IPSec		
IPsec Support for IPv6	6.X	Supported
IPsec Support for OSPF3	6.X	Supported
IPsec Support for RIPng	6.X	Supported
ADDITIONAL FEATURES		
MAC-Forced Forwarding (Dynamic Proxy ARP)	6.4.3	Not supported
Web Cache Communication Protocol (WCCP)	6.4.4	Not Supported
VDI Support	6.4.6	Not Supported
NTP Server	6.4.2	Supported
Traffic Anomaly Detection (Network Security)	6.X	Not Supported
UDLD	6.X	Supported
Router Discovery Protocol (RDP)	6.X	Not Supported
MAC Retention	6.X	Not Supported

Feature	Introduced	8.X
Interswitch Protocols (AMAP)	6.X	Not Supported
Pause Control/Flow Control	6.X	Supported
Autoboot Interruption	6.4.6	Not Supported