Release Notes

OmniSwitch 6900

Release 7.2.1.R01

These release notes accompany release 7.2.1.R01 software. 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 7 User Guides. The following are the titles and descriptions of the user manuals that apply to this release. User manuals can be downloaded at: http://enterprise.alcatel-lucent.com/?dept=UserGuides&page=Portal

- **OmniSwitch 6900 Series Getting Started Guide**
  Describes the hardware and software procedures for getting an OmniSwitch 6900 Series switch up and running.
- **OmniSwitch 6900 Series Hardware User Guide**
  Complete technical specifications and procedures for all OmniSwitch 6900 Series chassis, power supplies, and fans.
- **OmniSwitch CLI Reference Guide**
  Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.
- **OmniSwitch AOS Release 7 Network Configuration Guide**
  Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethe rnet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.
- **OmniSwitch AOS Release 7 Switch Management Guide**
  Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).
- **OmniSwitch AOS Release 7 Advanced Routing Configuration Guide**
  Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, and OSPFv3.
- **OmniSwitch AOS Release 7 Transceivers Guide**
  Includes SFP and SFP+ transceiver specifications and product compatibility information.
- **Technical Tips, Field Notices**
  Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

- OmniSwitch 6900 Series Release 7.2.1.R01 requires 2GB of SDRAM and 2GB flash memory. This is the standard configuration shipped.
- Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory.

UBoot and FPGA Requirements

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

Switches not running the minimum version required should upgrade to the latest Uboot, Miniboot, FPGA that is available with the 7.2.1.R01 AOS software available from Service & Support.

OmniSwitch 6900

Release	UBoot	FPGA/CPLD CMM	FPGA/CPLD Expansion Module
7.2.1.R01 (GA)	7.2.1.R01.117	1.0.0	0.2.0

```
-> show slot

Module in slot 1
 Model Name:                     OS6900-X40,
 Module Type:                    0x5062202,
 Description:                    40 SFP+ W/2EXP SLOT
 Part Number:                    055535-48T,
 Hardware Revision:              B,
 Serial Number:                  M0960002,
 Manufacture Date:               Mar  2 2011,
 FPGA/CPLD - Physical 1:         1.0.0,
 FPGA/CPLD - Physical 2:         1.0.0,
 Admin Status:                   POWER ON,
 Operational Status:             UP,
 Max Power:                      230,
 CPU Model Type:                 MPC 8572,
 MAC Address:                    00:e0:b1:e7:09:aa,
 ASIC - Physical 1:              BCM56845_B0,
 UBOOT Version:                  7.2.1.R01.117

Module in slot 2
 Model Name:                     OS-XNI-U4,
 Module Type:                    0x5072204,
 Description:                    4 SFP+,
 Part Number:                    050544-28T,
 Hardware Revision:              B02,
 Serial Number:                  M2460060,
 Manufacture Date:               Jun 16 2011,
 FPGA/CPLD - Physical 1:         0.2.0,
 Admin Status:                   POWER ON,
 Operational Status:             UP,
 Max Power:                      15,
 CPU Model Type:                 N/A,
 MAC Address:                    e8:e7:32:07:a0:a0,
 ASIC - Physical 1:              N/A,
 UBOOT Version:                  7.2.1.R01.117
```

Hardware Supported

OmniSwitch 6900-X20

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

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

OmniSwitch 6900-X40

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

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

OS-XNI-U4

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

OS-XNI-U12

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

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

450W modular AC power supply with front to back cooling.

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

450W modular DC power supply with front to back cooling.

OS6900-FT-F FanTray

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

Software Features Supported

The following software features are supported with the 7.2.1 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 an Advanced license.

Feature Summary

Feature	Platform	License
Manageability Feature Support		
CLI	OS6900	Base
Ethernet Interfaces	OS6900	Base
License Management	OS6900	Base
Multiple VRF Routing and Forwarding	OS6900	Advanced
Network Time Protocol (NTP) Client	OS6900	Base
Pause Control(RX) /Flow Control	OS6900	Base
Remote Access • FTP • SCP • SSH/SFTP • Telnet • TFTP	OS6900	Base
Resiliency Features • Hot Swap Expansion Modules • Power Supply Redundancy • Fan Redundancy	OS6900	Base
SNMP	OS6900	Base
Software Rollback – Multi-Image/Multi-Config	OS6900	Base
Storm Control	OS6900	Base
Text File Configuration	OS6900	Base
UDLD	OS6900	Base
USB Support	OS6900	Base
Web-Based Management (WebView)	OS6900	Base
Layer 2 Feature Support		
802.1AB with MED Extensions	OS6900	Base
802.1Q	OS6900	Base
Configurable Hash Mode	OS6900	Base
HA-VLAN	OS6900	Base
Link Aggregation –Static and LACP (802.3ad)	OS6900	Base
Multi-Chassis Link Aggregation	OS6900	Base
MVRP	OS6900	Base
Source Learning	OS6900	Base
Spanning Tree	OS6900	Base

Feature	Platform	License
• 802.1d and 802.1w • Multiple Spanning Tree Protocol • PVST+ • Root Guard		
Universal Network Profiles (UNP) – Chris	OS6900	Base
VLANs	OS6900	Base
IPv4 Feature Support		
Bi-Directional Forwarding Detection (BFD)	OS6900	Base
DHCP / UDP • DHCP Relay/Option-82 • Per-VLAN • UDP Relay	OS6900	Base
BGP4 with Graceful Restart	OS6900	Advanced
DNS Client	OS6900	Base
GRE	OS6900	Base
IP Multicast Routing	OS6900	Advanced
IP Multicast Switching (IGMP)	OS6900	Base
IP Multicast Switching (Proxying)	OS6900	Base
IP Multinetting	OS6900	Base
IP Route Map Redistribution	OS6900	Base
IP-IP Tunneling	OS6900	Base
OSPFv2	OS6900	Advanced
RIPv1/v2	OS6900	Base
Routing Protocol Preference	OS6900	Base
Server Load Balancing	OS6900	Base
VRRPv2	OS6900	Advanced
IPv6 Feature Support		
BGP4 • BGP IPv6 Extensions	OS6900	Advanced
IPSec IPv6 • OSPFv3 • RIPng	OS6900	Advanced
IPv6 Client and/or Server Support	OS6900	Base
IPv6 Multicast Routing	OS6900	Advanced
IPv6 Multicast Switching (MLD v1/v2)	OS6900	Base
IPv6 Routing	OS6900	Advanced
IPv6 Scoped Multicast Addresses	OS6900	Base
IPv6 Neighbor Discovery Support	OS6900	Base
OSPFv3	OS6900	Advanced
RIPng	OS6900	Advanced
VRRPv3	OS6900	Advanced
QoS Feature Support		

Feature	Platform	License
Auto-Qos Prioritization of NMS Traffic	OS6900	Base
Ingress and egress bandwith shaping	OS6900	Base
Policy Based Routing	OS6900	Advanced
Tri-Color Marking	OS6900	Base
Multicast Feature Support		
DVMRP	OS6900	Advanced
IGMP Multicast Group Configuration Limit	OS6900	Base
IGMP Relay	OS6900	Base
IPv4/IPv6 Multicast Switching (IPMS)	OS6900	Base
L2 Static Multicast Address	OS6900	Base
PIM / PIM-SSM (Source-Specific Multicast)	OS6900	Advanced
Monitoring/Troubleshooting Feature Support		
DDM - Digital Diagnostic Monitoring	OS6900	Base
Health Statistics	OS6900	Base
Ping and Traceroute	OS6900	Base
Policy Based Mirroring	OS6900	Base
Port Mirroring	OS6900	Base
Port Monitoring	OS6900	Base
Remote Port Mirroring	OS6900	Base
Rmon	OS6900	Base
sFlow	OS6900	Base
Switch Logging and Syslog	OS6900	Base
Metro Ethernet Feature Support		
ERP G.8032 – Shared VLAN	OS6900	Base
Ethernet Services	OS6900	Base
L2 Control Protocol Tunneling (L2CP)	OS6900	Base
Security Feature Support		
Access Control Lists (ACLs) for IPv4/IPv6	OS6900	Base
Account & Password Policies	OS6900	Base
Admin User Remote Access Restriction Control	OS6900	Base
ARP Defense Optimization	OS6900	Base
ARP Poisoning Detect	OS6900	Base
Authenticated Switch Access	OS6900	Base
IP DoS Filtering	OS6900	Base
Learned Port Security (LPS)	OS6900	Base
Policy Server Management	OS6900	Base

Manageability Feature Support

Command Line Interface (CLI)

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

Ethernet Interfaces

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

License Management

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

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

Advanced License Features

Multiple Virtual Routing and Forwarding (Multiple-VRF)

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

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

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

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

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

The Multiple VRF feature uses a context-based command line interface (CLI). When the switch boots up, a default VRF instance is automatically created and active. Any commands subsequently entered apply to this default instance. If a different VRF instance is selected, then all subsequent commands apply to that

instance. The CLI command prompt indicates which instance is the active VRF CLI context by adding the name of the VRF instance as a prefix to the command prompt (for example, **`vrf1: ->`**).

VRF - QoS

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

VRF - Switch Authentication

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

VRF - Switch Access and Utilities

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

VRF - VRRP

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

VRF – UDP/DHCP Relay

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

- VRF-Aware commands. These commands are allowed in any of the VRF instances configured in the switch. The settings in one VRF are independent of the settings in another VRF. Command parameters are visible and configurable within the context of any VRF.
- Global commands. These commands are supported only in the default VRF, but are visible and applied to all VRF instances configured in the switch. This command behavior is similar to how command parameters are applied in the per-VLAN DHCP Relay mode. For example, the maximum hops value configured in the default VRF is applied to all DHCP Relay agents across all VRF instances. This value is not configurable in any other VRF instance.
- Default VRF commands. These commands are supported only in the default VRF and are not applied to any other VRF instance configured in the switch. For example, per-VLAN mode and boot-up commands fall into this category.

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

VRF – PIM and DVMVRP

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

Network Time Protocol (NTP) Client

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.

Pause Control (RX)/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 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)**
  SCP is used secure manner between hosts on the network. The scp utility performs encrypted data transfers using the Secure Shell (SSH) protocol. In addition, scp uses available SSH authentication and security features, such as prompting for a password if one is required.
- **Secure Shell (SSH)/Secure FTP (SFTP)**
  The Secure Shell feature provides a secure mechanism that allows you to log in to a remote switch, to execute commands on a remote device, and to move files from one device to another. Secure Shell provides secure, encrypted communications even when your transmission is between two untrusted hosts or over an 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.

Hardware Resiliency

All OmniSwitch 6900 models support 1+1 redundant, hot-swappable AC and DC power supplies. The primary and backup power supply units are internal, but removable allowing for easier maintenance and replacement.There is no interruption of service when a new power supply is installed or an old one replaced. Additionaly the switch supports hot-swapping of the fan tray and plug-in modules of the same type.

SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that allows communication between SNMP managers and SNMP agents on an IP network. Network administrators use SNMP to monitor network performance and to solve network problems. SNMP provides an industry standard communications

model used by network administrators to manage and monitor their network devices. The OmniSwitch supports SNMPv1, SNMPv2, and SNMPv3.

Software Rollback – Multi-image/Multi-Config

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

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

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

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

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

Storm Control

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

Text File Configuration

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

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

UDLD - Fiber and Copper

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

USB Support

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

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

Web-Based Management (WebView)

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

- IE6, IE7, IE8 for Windows XP
- IE8, Firefox Mozilla 3.6 on Vista

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

Layer 2 Feature Support

802.1AB MED Extensions

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

802.1Q

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

Configurable Hash Mode

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

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

Default Hashing Mode and Recommendations

Platform	Default Hashing Mode
OS6900	Brief

- Changing the hash mode affects all features that rely on hashing, including Link Aggregation, ECMP and Server Load Balancing. Changing the hash mode per feature is not supported.

- Server Load Balancing uses dynamic port assignment, therefore it is not recommended to enable the TCP/UDP port hashing option with extended mode when SLB is configured on the switch.

High Availability -VLAN

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

HA VLAN Operational Modes

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

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

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

Link Aggregation - Static & LACP (802.3ad)

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

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

Non-Unicast Load Balancing on Link Aggregation

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

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

Multi-Chassis Link Aggregation

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

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

MVRP - Multiple VLAN Registration Protocol

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

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

Source Learning

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

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

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

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

Spanning Tree

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

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

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

Multiple Spanning Tree Protocol (MSTP)

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

PVST+ Interoperability

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

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

Universal Network Profile (UNP)

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

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

A UNP consists of the following attributes:

- **UNP Name**. The UNP name is obtained from the RADIUS server and mapped to the same profile name configured on the switch. The switch profile then identifies three attribute values: VLAN ID, Host Integrity Check (HIC) status, and a QoS policy list name.
- **VLAN ID**. All members of the profile group are assigned to the VLAN ID specified by the profile.
- **QoS Policy List Name**. Specifies the name of an existing list of QoS policy rules. The rules within the list are applied to all members of the profile group to enforce access to network resources. Only one policy list is allowed per profile, but multiple profiles may use the same policy list.

VLANs

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

- Creating or modifying VLANs.
- Assigning or changing default VLAN port associations (VPAs).
- Enabling or disabling VLAN participation in the current Spanning Tree algorithm.
- Enabling or disabling classification of mobile port traffic by 802.1Q tagged VLAN ID.
- Enabling or disabling VLAN authentication.

- Enabling or disabling unique MAC address assignments for each router VLAN defined.
- Displaying VLAN configuration information.

IPV4 Feature Support

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

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

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

BGP4

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

BGP Graceful Restart

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

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.

DHCP / UDP Relay

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

DHCP Relay Agent Information Option-82

The DHCP Option-82 feature enables the relay agent to insert identifying information into client-originated DHCP packets before the packets are forwarded to the DHCP server.

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

Per-VLAN DHCP Relay

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

UDP Relay

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

DNS Client

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

Generic Routing Encapsulation

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

IP Multinetting

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

IP Route Map Redistribution

Route map redistribution provides the ability to control which routes from a source protocol are learned and distributed into the network of a destination protocol. A route map consists of one or more user- defined

statements that can determine which routes are allowed or denied access to the network. In addition, a route map may also contain statements that modify route parameters before they are redistributed.

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

IP-IP Tunneling

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

OSPFv2

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

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

RIPv1/RIPv2

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

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

RIP Timer Configuration

- Update —The time interval between advertisement intervals.
- Invalid—The amount of time before an active route expires and transitions to the garbage state.
- Garbage—The amount of time an expired route remains in the garbage state before it is removed from the RIB.
- Holddown—The amount of time during which a route remains in the hold-down state.

Routing Protocol Preference

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

Server Load Balancing (SLB)

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

Using SLB clusters can provide cost savings (costly hardware upgrades can be delayed or avoided), scal-ability (as the demands on your server farm grow you can add additional physical servers), reliability (if one

physical server goes down the remaining servers can handle the remaining workload), and flexibility (you can tailor workload requirements individually to servers within a cluster).

Server Load Balancing - WRR

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

VRRPv2

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

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

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

IPV6 Feature Support

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

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

Globally Unique Local Unicast Addresses

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

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

Scoped Multicast Addresses

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

BGP4

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

BGP IPv6 Extensions

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

IPsec Support for IPv6, OSPFv3, RIPng

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

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

OSPFv3

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

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

RIPng

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

VRRPv2/VRRPv3

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

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

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

QoS Feature Support

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

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

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

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

Auto-Qos Prioritization for NMS Traffic

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

Ingress and Egress Bandwidth Shaping

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

Policy Based Routing (Permanent Mode)

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

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

Policy Lists

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

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

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

Redirect Policies (Port and Link Aggregate)

Two policy action commands are available for configuring QoS redirection policies: policy action redirect port and policy action redirect linkagg. A redirection policy sends traffic that matches the policy to a specific port or link aggregate instead of the originally intended destination. This type of policy may use any condition; the policy action determines which port or link aggregate to which the traffic is sent.

Tri-Color Marking

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

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

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

There are two types of TCM marking supported:

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

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

Multicast Feature Support

DVMRP

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

IGMP Multicast Group Configuration Limit

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

IGMP Relay - Relay IGMP Packets to Specific Host

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

IP Multicast Switching (IPMS) – IPv4/IPv6

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

Destination hosts signal their intent to receive a specific multicast stream by sending a request to do so to a nearby switch using Internet Group Management Protocol (IGMP). The switch then learns on which ports multicast group subscribers are attached and can intelligently deliver traffic only to the respective ports. This mechanism is often referred to as IGMP snooping (or IGMP 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.

L2 Static Multicast Addresses

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

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

PIM-SM/PIM-DM/PIM-SSM

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

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

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

Monitoring and Troubleshooting Feature Support

DDM - Digital Diagnostic Monitoring

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

- Temperature
- Supply Voltage
- Current
- Output (Transmit) Power
- Input (Receive) Power

Traps can be enabled if any of these above values crosses the pre-defined low or high thresholds of the transceiver.

Note: Not all transceivers support DDM, refer to the Transceivers Guide for additional DDM information.

Health Statistics

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

- Switch-level input/output, memory and CPU utilization levels
- Module-level and port-level input/output utilization levels
- For each monitored resource, the following variables are defined:
- Most recent utilization level (percentage)
- Average utilization level over the last minute (percentage)
- Average utilization level over the last hour (percentage)
- Maximum utilization level over the last hour (percentage)
- Threshold level

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

Ping and Traceroute

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

Port Mirroring

Port mirroring allows transmitted and received traffic from a "mirrored" port to be copied to another port. The "mirroring" port receives a copy of all 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 limitations:

- The policy mirror action must specify the same analyzer port for all policies in which the action is used.
- One policy-based mirroring session supported per switch.
- One port-based mirroring session supported per switch. Note that policy-based and port-base mirroring are both allowed on the same port at the same time.
- One remote port-based mirroring session supported per switch.
- One port-monitoring session supported per switch.

Port Mirroring – Remote (802.1Q Based)

This feature provides a remote port mirroring capability where traffic from a local port can be carried across the network to an egress port where a sniffer can be attached. This 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.

Metro Ethernet Feature Support

Ethernet Ring Protection (ERP) – G.8032

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

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

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

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

Ethernet Services

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

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

- Ethernet service-based approach that is similar to configuring a virtual private LAN service (VPLS).
- Ingress bandwidth sharing across User Network Interface (UNI) ports.
- Ingress bandwidth rate limiting on a per UNI port, per CVLAN, or CVLAN per UNI port basis.
- CVLAN (inner) tag 802.1p-bit mapping to SVLAN (outer) tag 802.1p bit.
- CVLAN (inner) tag DSCP mapping to SVLAN (outer) tag 802.1p bit.
- Profiles for saving and applying traffic engineering parameter values.
- Capability to suspend the use of SAP bandwidth and priority actions allowing QoS rules for advanced classification of SAP traffic, such as mapping several DSCP/ToS values to the same outer 802.1p value.

Ethernet Services - Egress Rate Limiting

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

- Destination port = UNI port defined in the sap
- VLAN = CVLAN defined in the sap (could be untagged, cvlan all or specific vlan id)
- Rate limiter with the sap-profile egress-bandwidth

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

Ethernet Services - Tunneling L2 Protocols

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

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

Security Feature Support

Access Control Lists (ACLs)

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

- **Layer 2 ACLs**—for filtering traffic at the MAC layer. Usually uses MAC addresses or MAC groups for filtering.
- **Layer 3/4 ACLs**—for filtering traffic at the network layer. Typically uses IP addresses or IP ports for filtering; note that IPX filtering is not supported.
- **Multicast ACLs**—for filtering IGMP traffic.
- **ICMP drop rules**—Allows condition combinations in policies that will prevent user pings, thus reducing DoS exposure from pings. Two condition parameters are also available to provide more granular filtering of ICMP packets: icmptype and icmpcode.
- **TCP connection rules**—Allows the determination of an established TCP connection by examining TCP flags found in the TCP header of the packet. Two condition parameters are available for defining a TCP connection ACL: established and tcpflags.
- **Early ARP discard**—ARP packets destined for other hosts are discarded to reduce processing overhead and exposure to ARP DoS attacks. No configuration is required to use this feature, it is always available and active on the switch. Note that ARPs intended for use by a local subnet, AVLAN, and VRRP are not discarded.
- **UserPorts**—A port group that identifies its members as user ports to prevent spoofed IP traffic. When a port is configured as a member of this group, packets received on the port are dropped if they contain a source IP network address that does not match the IP subnet for the port.
- **UserPorts Profile**—In addition to spoofed traffic, it is also possible to configure a global UserPorts profile to specify additional types of traffic, such as BPDU, RIP, OSPF, DVMRP, PIM, DHCP server response packets, DNS and/or BGP, to monitor on user ports. The UserPorts profile also determines whether user ports will filter the unwanted traffic or will administratively shutdown when the traffic is received. Note that this profile only applies to those ports that are designated as members of the UserPorts port group.
- **DropServices**—A service group that improves the performance of ACLs that are intended to deny packets destined for specific TCP/UDP ports. This group only applies to ports that are members of the UserPorts group. Using the DropServices group for this function minimizes processing overhead, which otherwise could lead to a DoS condition for other applications trying to use the switch.

Access Control Lists (ACLs) for IPv6

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

- Trusted/untrusted behavior is the same for IPv6 traffic as it is for IPv4 traffic.
- IPv6 policies do not support the use of network groups, service groups, map groups, or MAC groups.
- IPv6 multicast policies are not supported.
- Anti-spoofing and other UserPorts profiles/filters do not support IPv6.
- The default (built-in) network group, “Switch”, only applies to IPv4 interfaces. There is no such group for IPv6 interfaces.

Account and Password Policies

This feature allows a switch administrator to configure password policies for password creation and management. The administator 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 entires to the forwarding table until the net hop ARP entry can be resolved.

ARP Poisoning Detect

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

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

Authenticated Switch Access

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

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

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

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

IP DoS Filtering

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

- ARP Flood Attack
- Invalid IP Attack
- Multicast IP and MAC Address Mismatch
- Ping Overload
- Packets with loopback source IP address

Learned Port Security (LPS)

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

- A configurable source learning time limit that applies to all LPS ports.
- A configurable limit on the number of MAC addresses allowed on an LPS port.
- Dynamic configuration of a list of authorized source MAC addresses.
- Static configuration of a list of authorized source MAC addresses.
- Two methods for handling unauthorized traffic: Shutting down the port or only blocking traffic that violates LPS criteria.
- A configurable limit to the number of filtered MAC addresses allowed on an LPS port. Conversion of dynamically learned MAC addresses to static MAC address entries.
- Support for all authentication methods and LPS on the same switch port.

LPS has the following limitations:

- You cannot configure LPS on link aggregate ports.

Learned MAC Address Notification

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

SNMP Traps

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

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

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

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

No.	Trap Name	Platforms	Description
			MAC.
54	gvrpVlanLimitReachedEvent	OS6900	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
55	alaNetSecPortTrapAnomaly	OS6900	Trap for an anomaly detected on a port.
56	alaNetSecPortTrapQuarantine	OS6900	Trap for an anomalous port quarantine.
57	ifMauJabberTrap		This trap is sent whenever a managed interface MAU enters the jabber state.
58	udldStateChange	OS6900	Generated when the state of the UDLD protocol changes.
59	ndpMaxLimitReached	OS6900	This IPv6 Trap is sent when the hardware table has reached the maximum number of entries supported.
60	ripRouteMaxLimitReached	OS6900	This trap is sent when the RIP database reaches the supported maximum number of entries. When the maximum number is reached, RIP discards any new updates.
61	ripngRouteMaxLimitReached	OS6900	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	OS6900	This trap is sent when the ERP Ring State has changed from “Idle” to “Protection”.
63	alaErpRingMultipleRpl	OS6900	This trap is sent when multiple RPLs are detected in the Ring.
64	alaErpRingRemoved	OS6900	This trap is sent when the Ring is removed dynamically.
65	ntpMaxAssociation	OS6900	This trap is generated when the maximum number of peer and client associations config-ured for the switch is exceeded.
66	ddmTemperatureThresholdViolated	OS6900	This trap is sent when an SFP/ XFP/SFP+ temperature has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/ XFP/SFP+ temperature.
67	ddmVoltageThresholdViolated	OS6900	This trap is sent when SFP/XFP/ SFP+ supply voltage has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ supply voltage.
68	ddmCurrentThresholdViolated	OS6900	This trap is sent when if an SFP/ XFP/SFP+ Tx bias current has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx bias current.
69	ddmTxPowerThresholdViolated	OS6900	This trap is sent when an SFP/ XFP/SFP+ Tx output power has crossed any threshold or reverted from previous threshold violation for a

No.	Trap Name	Platforms	Description
			port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx output power.
70	ddmRxPowerThresholdViolated	OS6900	This trap is sent when an SFP/ XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
71	webMgtServerErrorTrap	OS6900	This trap is sent to management station(s) when the Web Management server goes into error state after becoming unreachable twice within a minute.
72	multiChassisIpcVlanUp	OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is up.
73	multiChassisIpcVlanDown	OS6900	This trap is sent to indicate the operational status for the multi-chassis communication VLAN is down.
74	multiChassisMisconfigurationFailure	OS6900	This trap is sent to indicate a mis-configuration due to Chassis Id or IPC VLAN.
75	multiChassisHelloIntervalConsisFailure	OS6900	This trap is sent to indicate a hello interval consistency failure.
76	multiChassisStpModeConsisFailure	OS6900	This trap is sent to indicate an STP mode consistency failure.
77	multiChassisStpPathCostModeConsisFailure	OS6900	This trap is sent to indicate an STP path cost mode consistency failure.
78	multiChassisVflinkStatusConsisFailure	OS6900	This trap is sent to indicate a VFLink status consistency failure.
79	multiChassisStpBlockingStatus	OS6900	This trap is sent to indicate the STP status for some VLANs on the VFLink is in blocking state.
80	multiChassisLoopDetected	OS6900	This trap is sent to indicate a loop has been detected.
81	multiChassisHelloTimeout	OS6900	This trap is sent to indicate the hello timeout has occurred.
82	multiChassisVflinkDown	OS6900	This trap is sent to indicate the VFLink is down.
83	multiChassisVFLMemberJoinFailure	OS6900	This trap is sent to indicate a port configured as virtual-fabric member is unable to join the virtual-fabric link.
84	alaDHLVlanMoveTrap	OS6900	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	OS6900	This trap is sent when a new IP address is assigned to DHCP Cli-ent interface.
86	alaDhcpClientAddressExpiryTrap	OS6900	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address.
87	alaDhcpClientAddressModifyTrap	OS6900	This trap is sent when the DHCP client is unable

No.	Trap Name	Platforms	Description
			to obtain the existing IP address and a new IP address is assigned to the DHCP client
88	vRtrIsisDatabaseOverload	OS6900	This notification is generated when the system enters or leaves the overload state.
89	vRtrIsisManualAddressDrops	OS6900	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
90	vRtrIsisCorruptedLSPDetected	OS6900	This notification is generated when an LSP that was stored in memory has become corrupted.
91	vRtrIsisMaxSeqExceedAttempt	OS6900	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
92	vRtrIsisIDLenMismatch	OS6900	A notification sent when a PDU is received with a different value of the System ID Length.
93	vRtrIsisMaxAreaAddrsMismatch	OS6900	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
94	vRtrIsisOwnLSPPurge	OS6900	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
95	vRtrIsisSequenceNumberSkip	OS6900	When an LSP is received without a System ID and different contents.
96	vRtrIsisAutTypeFail	OS6900	A notification sent when a PDU is received with the wrong authentication type field.
97	vRtrIsisAuthFail	OS6900	A notification sent when a PDU is received with an incorrent authentication information field.
98	vRtrIsisVersionSkew	OS6900	A notification sent when a a Hello PDU is received from an IS running a different version of the protocol.
99	vRtrIsisAreaMismatch	OS6900	A notification sent when a Hello PDU is received from an IS which does not share any area address.
100	vRtrIsisRejectedAdjacency	OS6900	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	OS6900	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
102	vRtrIsisOrigLSPBufSizeMismatch	OS6900	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	OS6900	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
104	vRtrIsisAdjacencyChange	OS6900	A notification sent when an adjacency changes state, entering or leaving state up. The first 6

No.	Trap Name	Platforms	Description
			bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
105	vRtrIsisCircIdExhausted	OS6900	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	OS6900	A notification sent when an adjancency's graceful restart status changes.
107	mvrpVlanLimitReachedEvent	OS6900	This trap is sent when the num-ber of VLANs learned dynami-cally by MVRP reaches the configured limit.

Unsupported Software Features

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

Feature	Platform	License
Dual-Home Link Aggregation	OS6900	Base
IS-IS	OS6900	Advanced
NetSec	OS6900	Base

Unsupported CLI Commands

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

Software Feature	Unsupported CLI Commands
Qos	qos wrp qos qsp show qos qsi wred-stats
Source Learning	mac-learning mode [distributed | centralized]

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.

Layer 2

PR	Description	Workaround
95308	Temporary traffic loops could happen under the following scenarios: 1. Reloading of a non root bridge. This happens when the bridge is going down and is due to the sequential bringing down of NIs during a reload process .It is purely temporary in nature and stops when all the NIs eventually get powered off. 2. NI power down When an NI power down command is executed for an NI and if that NI has the Root port port and other NIs have Alternate ports, it is possible to see some traffic looping back from the newly elected Root port. The traffic loop back is temporary and will stop once the NI gets powered off. 3. New Root bridge selection Temporary loops could occur during the process of electing a new Root bridge, if this election process is triggered by the assignment a worse priority for the existing root bridge or a root bridge failure. This happens due to the inconsistent spanning tree topology during the convergence and stops entirely once the network converges	For items 1 and 2 above there is no work around presently. For item 3 the following work around could be applied: 1. Tune the max age (and or max hops in the case of MSTP) parameter to a lower value that is optimal for the network. This will reduce the convergence time and thereby the duration of temporary loops. 2. To select a new root bridge, consider assigning better priority for that bridge instead of assigning worse priority for the existing root bridge.
159973	Auto-negotiation does not remain disabled after a reboot if auto-negotiation was disabled on a port detected as 10-Gigabit and the switch is rebooted with a 1-Gigabit transceiver.	Ensure the port is detected as 1-Gigabit before disabling auto-negotiation and rebooting.
160462	Disabling MAC learning for a VLAN containing a UNP port can cause UNP to stop classifying traffic.	Do not disable mac-learning for a VLAN containing UNP ports.
160705	During times of heavy source learning activity the command 'show mac-learning summary vlan <vlan>' may cause "Please wait ..." to be displayed on the console.	There is no known workaround at this time.

Layer 3

PR	Description	Workaround
159042	OS6900 supports ARP learning at a maximum rate of 2000 pps. At higher rates, some ARPs may not be learned.	There is no known workaround at this time.
159145	By default, packets that originate from a switch with an address assigned to a VRRP virtual router will use the real hardware MAC as their source, not the VRRP virtual MAC. The 'ipv6 virtual-source-mac on|off' command can be used to modify this behaviour. If 'on' is specified, the VRRP virtual MAC will be used as the source. This command has no effect on VRRP advertisements, which will always be sent using the VRRP virtual MAC as the source.	Use the 'ipv6 virtual-source-mac on|off' command to change the behavior as desired.
159248	When more than 256 IPv6 routes with a prefix length greater than 64 bits are active, some traffic will be routed in software.	There is no known workaround at this time.
160068	For each bootp request received, two are routed out (one is correct, the other has the wrong MAC address.	There is no known workaround at this time.
160746	A VRF mismatch message similar to "bcd itf alarm message: >> +++ vrf mismatch 143 (0 0x11)" may be seen when trying to create an IP interface.	When an IP interface (v4 or v6) is bound to a VLAN that VLAN can only be used in the same vrf as the original interface. This error occurs when trying to create another IP interface for that VLAN in a different vrf.
160769	Wake-on-LAN packets are not forwarded by UDP relay.	There is no known workaround at this time.

Port Mirroring/Monitoring

PR	Description	Workaround
159281	Traffic is captured as it is on ingress, not as it should be on egress on a mirrored port (sometimes packets are modified prior to being sent out).	There is no known workaround at this time.
159319	Port monitoring only captures the first outgoing ICMP reply, the remaining are not captured.	There is no known workaround at this time.

QoS

PR	Description	Workaround
160678	The creation of non-default QoS profiles is allowed even though only the default profile is supported.	There is no known workaround at this time.
160458	After removing a policy which does not exist, an 'incorrect slot/port' error message may be displayed.	There is no known workaround at this time.

Security

PR	Description	Workaround
160557	If a MAC address is learned on a UNP port, the same MAC address cannot be learned on another UNP port on any other VLAN (tagged or untagged).	There is no known workaround at this time.
160766	AAA allows a user password size of 6 characters. However, if the same user is polled from OmniVista using SNMP, the switch will report an authentication failure because OmniVista is expecting a minimum password size of 8 characters.	Create a password with a minimum of 8 characters.
160818	UNP is not able to classify users into a dynamic VLAN created by MVRP.	There is no known workaround at this time.

System

PR	Description	Workaround
157986	A CLI session timeout can interrupt a file transfer to the USB flash drive, causing the transfer to fail.	When transferring large files, increase the CLI session timeout using the 'session cli timeout' command.

WebView

PR	Description	Workaround
153219	WebView does not display the switch log.	View the switch log files using the CLI using respective commands (more, vi, etc.).

Hot Swap/Redundancy Feature Guidelines

Hot Swap Feature Guidelines

- Hot swap of like modules is supported.
- Hot swap of unlike modules is not supported.
- Hot insertion, the insertion of a module into a previously empty slot, is supported.

Hot Swap Procedure

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

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

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

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

Technical Support

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

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

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

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

Worldwide (except Noth America) web: https://businessportal.alcatel-lucent.com

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

Third Party Licenses and Notices

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

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