

Release Notes

OmniSwitch 6800/6850/9000

Release 6.3.1.R01

These release notes accompany release 6.3.1.R01 software for the OmniSwitch 6800, 6850, and 9000 hardware. They 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**, see page 3.
- **System Requirements**, see page 5
 - Memory Requirements
 - UBoot, FPGA, MiniBoot, BootROM, and Upgrade Requirements
- **New Hardware Supported**, see page 7
- **Supported Hardware/Software Combinations**, see page 8
- **New Software Features**, see page 10
- **Software Supported**, see page 11
- **Supported Traps**, see page 45
- **Unsupported Software Features**, see page 52
- **Unsupported CLI Commands**, see page 52
- **Unsupported MIBs**, see page 54
- **Open Problem Reports, and Feature Exceptions**, see page 59
- **Technical Support**, see page 71

Related Documentation

These release notes should be used in conjunction with the OmniSwitch 6800, 6850, and 9000. The following are the titles and descriptions of the user manuals that apply to the OmniSwitch 6800, 6850, and 9000.

User manuals can be downloaded at:

http://www1.alcatel-lucent.com/enterprise/en/resource_library/user_manuals.html.

- ***OmniSwitch 6800 Series Getting Started guide***
Describes the hardware and software procedures for getting an OmniSwitch 6800 Series switch up and running.
- ***OmniSwitch 6850 Series Getting Started Guide***
Describes the hardware and software procedures for getting an OmniSwitch 6850 Series switch up and running.
- ***OmniSwitch 9000 Series Getting Started Guide***
Describes the hardware and software procedures for getting an OmniSwitch 9000 Series switch up and running.
- ***OmniSwitch 6800 Series Hardware User Guide***
Complete technical specifications and procedures for all OmniSwitch 6800 Series chassis, power supplies, and fans.
- ***OmniSwitch 6850 Series Hardware User Guide***
Complete technical specifications and procedures for all OmniSwitch 6850 Series chassis, power supplies, and fans.
- ***OmniSwitch 9000 Series Hardware User Guide***
Complete technical specifications and procedures for all OmniSwitch 9000 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 6800/6850/9000 Network Configuration Guide***
Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethernet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.
- ***OmniSwitch 6800/6850/9000 Series 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 6800/6850/9000 Series Advanced Routing Configuration Guide***

December 2007

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.

- ***Upgrade Instructions for 6.3.1.R01***

Provides instructions for upgrading the OmniSwitch 6800, 6850, 9000 to 6.3.1.R01.

- ***OmniSwitch Transceivers Guide***

Includes SFP and XFP 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 6800 Series Release 6.3.1.R01 requires 256 MB of SDRAM and 64MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 6850 Series Release 6.3.1.R01 requires 256 MB of SDRAM and 64MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 9000 Series Release 6.3.1.R01 requires 256 MB of SDRAM and 128MB of flash memory for the Chassis Management Module (CMM). 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. Use the **show hardware info** command to determine your SDRAM and flash memory.

UBoot, FPGA, Miniboot, BootROM, and Upgrade Requirements.

The software versions listed in this section are the minimum required, except where otherwise noted.

OmniSwitch 9000

Release	Miniboot.uboot CMM	UBoot CMM	UBoot NI	FPGA CMM
6.1.1.R01	-	6.1.1.632.R01	6.1.1.632.R01	2.20
6.1.1.R02	-	6.1.1.632.R01 (minimum) 6.1.1.109.R02 (recommended)	6.1.1.632.R01 (minimum) 6.1.1.109.R02 (recommended)	Major Revision: 2 Minor Revision: 23 (displays as 0x17; recommended)
6.1.3.R01	6.1.1.167.R02 or later	6.1.1.167.R02	6.1.1.167.R02	Major Revision: 2 Minor Revision: 25 (displays as 0x19; recommended)
6.1.5.R01	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	Major Revision: 2 Minor Revision: 25 (displays as 0x19; recommended)
6.3.1	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	6.1.1.167.R02 (minimum) 6.1.5.354.R01 recommended	Major Revision: 2 Minor Revision: 25 (displays as 0x19; recommended)

OmniSwitch 6850

Release	Miniboot.uboot	UBoot
6.1.2.R03	6.1.2.49.R03 (Miniboot)	6.1.3.601.R01 (BootROM)
6.1.3.R01	6.1.3.601.R01	6.1.3.601.R01
6.1.5.R01	6.1.3.601.R01	6.1.3.601.R01
6.3.1.R01	6.1.3.601.R01	6.1.3.601.R01

OmniSwitch 6800

Release	Miniboot	BootROM
6.1.2.R03	6.1.2.49.R03	6.1.2.49R03
6.1.3.R01	6.1.2.49.R03	6.1.2.49R03
6.1.5.R01	6.1.2.261.R03	6.1.2.261.R03
6.3.1	6.1.2.261.R03	6.1.2.261.R03

POE Firmware

- 5.01

Refer to the *Upgrading OmniSwitch 6800, 6850, and 9000 Series Switches to 6.3.1.R01* instructions if a switch upgrade is necessary to meet the above requirements

Note: Refer to the *Upgrading OmniSwitch 6800, 6850, and 9000 Series Switches to 6.3.1.R01* instructions if a switch upgrade is necessary to meet the above requirements.

New Hardware Supported

New Network Interface (NI) Modules

No new modules are introduced in this release.

SFPs

SFP-DUAL-MM (OS6850)

Dual Speed 100Base-FX or 1000Base-LX Ethernet optical transceiver. Supports multimode fiber and uses an LC connector. Release 6.3.1 provides support for this SFP on OS-6850 combo ports.

SFP-DUAL-SM10 (OS6850)

Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver. Supports single mode fiber and uses LC connector. Release 6.3.1 provides support for this SFP on OS-6850 combo ports.

SFP-Gig-BX-U (OS6850/OS9000)

1000Base-BX SFP transceiver with an LC type connector. This bi-directional transceiver is designed for use over single-mode fiber optic on a single strand link up to 10 km. Transmits at 1310 nm and receives at 1490 nm optical signal.

SFP-GIG-BX-D (OS6850/OS9000)

1000Base-BX SFP transceiver with an LC type connector. This bi-directional transceiver is designed for use over single-mode fiber optic on a single strand link up to 10 km. Transmits at 1490 nm and receives at 1310 nm optical signal.

Supported Hardware/Software Combinations

The following table shows the 6.X software releases that support each of the listed OS9000, OS6850, and OS6800 module types:

Module Type	Part Number	6.1.1.R01	6.1.2.R03	6.1.3.R01	6.1.5.R01	6.3.1.R01
OS96/9700 CMM, REV B	902369	supported	n/a	supported	supported	supported
OS96/9700 CMM, REV C	902444	supported	n/a	supported	supported	supported
OS9800 CMM	902492	not supported	n/a	supported	supported	supported
OS9-GNI-C24, ASIC A1	902367	supported	n/a	supported	supported	supported
OS9-GNI-U24, ASIC A1	902370	supported	n/a	supported	supported	supported
OS9-XNI-U2, ASIC A1	902379	supported	n/a	supported	supported	supported
OS9-GNI-C20L, ASIC B2	902434	not supported	n/a	not supported	supported	supported
OS9-GNI-C24, ASIC B2	902394	not supported	n/a	supported	supported	supported
OS9-GNI-C48T, ASIC B2	902507	not supported	n/a	not supported	supported	supported
OS9-GNI-U24, ASIC B2	902396	not supported	n/a	supported	supported	supported
OS9-XNI-U2, ASIC B2	902397	not supported	n/a	supported	supported	supported
OS9-GNI-P24, ASIC B2	902395	not supported	n/a	supported	supported	supported
OS9-XNI-U6, ASIC B2	902398	not supported	n/a	supported	supported	supported
OS6850-24	902457	n/a	supported	supported	supported	supported
OS6850-48	902495	n/a	supported	supported	supported	supported
OS6850-24X	902458	n/a	supported	supported	supported	supported
OS6850-48X	902462	n/a	supported	supported	supported	supported
OS6850-P24	902459	n/a	supported	supported	supported	supported
OS6850-P48	902463	n/a	supported	supported	supported	supported
OS6850-P24X	902460	n/a	supported	supported	supported	supported
OS6850-P48X	902464	n/a	supported	supported	supported	supported
OS6850-U24X	902418	n/a	not supported	supported	supported	supported
OS6850-24L	902487	n/a	not supported	supported	supported	supported
OS6850-48L	902489	n/a	not supported	supported	supported	supported
OS6850-P24L	902488	n/a	not supported	supported	supported	supported
OS6850-P48L	902490	n/a	not supported	supported	supported	supported
OS6800-24	902349	n/a	supported	supported	supported	supported
OS6800-48	902350	n/a	supported	supported	supported	supported
OS6800-24L	902377	n/a	supported	supported	supported	supported
OS6800-48L	902378	n/a	supported	supported	supported	supported
OS6800-U24	902351	n/a	supported	supported	supported	supported

To determine the ASIC revision for a specific NI, use the **show ni** command. For example, the following **show ni** output display shows a B2 revision level for NI 1:

```
DC-Core ->> show ni 1
Module in slot 1
  Model Name:                OS9-GNI-C24,
  Description:               10-1000 RJ45,
  Part Number:               902394-40,
  Hardware Revision:         C13,
  Serial Number:             G1511279,
  Manufacture Date:          MAY 03 2006,
  Firmware Version:          ,
  Admin Status:              POWER ON,
  Operational Status:        UP,
  Power Consumption:          51,
  Power Control Checksum:    0x0,
  MAC Address:               00:d0:95:e6:54:80,
  ASIC - Physical 1:      BCM56504_B2
  CPLD - Physical 1:         0005/00
  UBOOT Version :            6.1.1.167.R02
  UBOOT-miniboot Version :   No Miniboot
  POE SW Version :           n/a
```

To determine the CMM board revision, use the **show cmm** command. For example, the following **show cmm** output display shows a C revision level for the CMM board:

```
DC-Core ->> show cmm
Module in slot CMM-A-1
  Model Name:                OS9700-CFM,
  Description:               FABRIC BOARD,
  Part Number:               902444-10,
  Hardware Revision:      C11,
  Serial Number:             G1810128,
  Manufacture Date:          MAY 08 2006,
  Firmware Version:          2,
  Admin Status:              POWER ON,
  Operational Status:        UP,
  Power Consumption:          27,
  Power Control Checksum:    0x0,
  MAC Address:               00:d0:95:e0:6c:ac,
```

New Software Features and Enhancements

The following software features and enhancements are new with the 6.3.1.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

Feature/Enhancement Summary

Feature	Platform	Software Package
31-bit Network Mask Support	OS6850/OS9000	base
802.1ab	all	base
Account & Password Policies	OS6850/OS9000	base
Quarantine Manager and Remediation	OS6800/OS6850	base
ARP Defense Optimization	OS6850/OS9000	base
ARP Poisoning Detect	all	base
Auto-Qos Prioritization of IP Phone Traffic	all	base
Auto-Qos Prioritization of NMS Traffic	all	base
AVLAN support for IE7/Windows Vista	all	base
DHCP Snooping Option-82 Data Insertion Format	all	base
DSCP Range Condition	OS6850/OS9000	base
ECMP RIP Support	all	base
Ethernet OAM	OS6850/OS9000	base
Generic Routing Encapsulation	OS6850/OS9000	base
GVRP	OS6850/OS9000	base
IP-IP Tunneling	OS6850/OS9000	base
IP MC VLAN – Support for multiple sender ports	OS6850/OS9000	base
IPv6 Client and/or Server Support	OS6850/OS9000	base
IPv6 Multicast Routing	OS6850/OS9000	advanced routing
IS-IS	OS6850/OS9000	advanced routing
Learned MAC Address Notificaton	all	base
L4 ACLs over IPv6	OS6850/OS9000	base
Mac Authentication for Supplicant/non-Supplicant	OS6850/OS9000	base
MAC Retention	OS6850	base
Traffic Anomaly Detection (Network Security)	OS6850/OS9000	base
Policy Based Mirroring	OS6850/OS9000	base
Port-based Ingress Limiting	OS6850/OS9000	base
PVST+	OS6850/OS9000	base
Remote Port Mirroring	OS6850/OS9000	base
RRSTP	OS9000	base
UDLD	OS6850/OS9000	base
User Network Profiles	All	base
VRRP Global Commands	OS6850/OS9000	base

VLAN Stacking Eservices	OS6850/OS9000	base
Webview/SNMP support for BGP IPv6 Extensions	all	advanced routing
Windows Vista for WebView	all	base

Software Supported

In addition to the new software features introduced with the 6.3.1.R01 release, the following software features are also supported in 6.3.1.R01, subject to the feature exceptions and problem reports described later in these release notes:

Feature Summary

Feature	Platform	Software Package
802.1Q	All	base
802.1Q 2005 (MSTP)	All	base
802.1x Multiple Client Support	All	base
802.1x Device Classification (Access Guardian)	all	base
Access Control Lists (ACLs)	all	base
Access Control Lists (ACLs) for IPv6	OS6850/OS9000	base
ACL & Layer 3 Security	all	base
ACL Manager (ACLMAN)	all	base
Authenticated Switch Access	all	base
Authenticated VLANs	all	base
Automatic VLAN Containment (AVC)	all	base
BGP4	all	base advanced routing
BGP Graceful Restart	all	base advanced routing
BPDU Shutdown Ports	OS6800	base
Command Line Interface (CLI)	all	base
DHCP Relay	all	base
DHCP Option-82	all	base
DHCP Snooping	all	base
DNS Client	all	base
Dynamic VLAN Assignment (Mobility)	all	base
DVMRP	all	base advanced routing
End User Partitioning	all	base
Ethernet Interfaces	all	base
Flood/Storm Control	all	base
Health Statistics	all	base
HTTP/HTTPS Port Configuration	all	base
Interswitch Protocols (AMAP)	all	base
IPv4 Routing	OS6800/OS6850/OS9000	base

IPv6 Routing	OS6850/OS9000	base
IP DoS Filtering	OS6850/OS9000	base
IPv4 Multicast Switching (IPMS)	all	base
IPv6 Multicast Switching (MLD)	OS6850/OS9000	base
IPv4 Multicast Switching (Proxying)	OS6800/OS6850/OS9000	base
IPv6 Multicast Switching (Proxying)	OS6850/OS9000	base
IP Multinetting	all	base
IP Route Map Redistribution	all	base
IPX Routing	all	base
L2 DHCP Snooping	all	base
L2 Static Multicast Address	all	base
Learned Port Security (LPS)	all	base
Link Aggregation (static & 802.3ad)	all	base
MAC Address Mode	OS9000	base
NTP Client	all	base
OSPFv2	all	base advanced routing
OSPFv3	OS6850/OS9000	base advanced routing
Partitioned Switch Management	all	base
Per-VLAN DHCP Relay	all	base
PIM	all	base
PIM-SSM (Source-Specific Multicast)		advanced routing
Policy Server Management	all	base
Policy Based Routing (Permanent Mode)	OS6850/OS9000	base
Port Mapping	all	base
Port Mirroring (1:24)	OS6800	base
Port Mirroring (1:128)	OS6850/OS9000	base
Port Monitoring	all	base
Power over Ethernet (PoE)	OS6850/OS9000	base
Quality of Service (QoS)	all	base
Redirection Policies (Port and Link Aggregate)	OS6850/OS9000	base
RIPv1/RIPv2	all	base
RIPng	OS6850/OS9000	base
RMON	all	base
Router Discovery Protocol (RDP)	all	base
Routing Protocol Preference	all	base
Secure Copy (SCP)	all	base
Secure Shell (SSH)	all	base
Server Load Balancing	OS6850/OS9000	base
SSH Public Key Authentication	all	base
sFlow	OS6850/OS9000	base
Smart Continuous Switching	all	base
Hot Swap		
Management Module Failover		
Power Monitoring		
Redundancy		

SNMP	all	base
Source Learning	all	base
Software Rollback	all	base
Spanning Tree	all	base
Syslog to Multiple Hosts	all	base
Switch Logging	all	base
Text File Configuration	all	base
User Definable Loopback Interface	all	base
VLANs	all	base
VLAN Stacking and Translation	OS6850/OS9000	base
VRRPv2	all	base
VRRPv3	OS6850/OS9000	base
Web-Based Management (WebView)	all	base

Feature Descriptions

31-Bit Network Mask Support

Adds support for a 31-bit netmask to allow for a point-to-point Ethernet network between two routers.

802.1AB with MED Extensions

IEEE 802.1AB (2005) is the latest version for the standards based connectivity discovery protocol. The purpose of the IEEE standard 802.1AB for Link Layer Discovery Protocol (LLDP), is to provide support for network management software dealing with topology discovery such as OmniVista. Switches that are compliant with 802.1AB exchange information with neighboring devices and maintain a database of the information exchanged. 802.1ab uses TLV (Time, Length, Value) frames to exchange information with neighboring devices. The **Link Layer Discovery Protocol-Media Endpoint Discover** or **LLDP-MED** is designed to extend IEEE 802.1AB functionality to exchange information such as VLANs and power capabilities.

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 area or identified as being destined for a specific area.

When a port is enabled to accept tagged traffic, by default both 802.1Q tagged and untagged traffic is automatically accepted on the port. Configuring the port to accept only tagged traffic is also supported.

802.1Q 2005 (MSTP)

802.1Q 2005 (Q2005) is a version of Multiple Spanning Tree Protocol (MSTP) that 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.

802.1x Device Classification (Access Guardian)

In addition to the authentication and VLAN classification of 802.1x clients (supplicants), this implementation of 802.1x secure port access extends this type of functionality to non-802.1x clients (non-supplicants). To this end device classification policies are introduced to handle both supplicant and non-suppliant access to 802.1x ports.

Supplicant policies use 802.1x authentication via a remote RADIUS server and provide alternative methods for classifying supplicants if the authentication process either fails or does not return a VLAN ID.

Non-suppliant policies use MAC authentication via a remote RADIUS server or can bypass authentication and only allow strict assignment to specific VLANs. MAC authentication verifies the source MAC address of a non-suppliant device via a remote RADIUS server. Similar to 802.1x

authentication, the switch sends RADIUS frames to the server with the source MAC address embedded in the username and password attributes.

In release 6.3.1 the number of possible 802.1X users is 2K per system, not to exceed 1K per module or stackable unit. This number is a total number of users that applies to all authenticated clients, such as AVLAN and 802.1X supplicants or non-supplicants. In addition the use of all authentication methods and Learned Port Security (LPS) on the same port is supported.

Release 6.3.1 added the capability to classify both supplicant and non-supplicant devices using non-supplicant device classification policies. As a result, MAC authentication is now applicable to both supplicant and non-supplicant devices.

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.

Access Control Lists (ACLs) for IPv6

Support for IPv6 ACLs on the OmniSwitch 6850 Series and OmniSwitch 9000 Series is available. The following QoS policy conditions are now available for configuring ACLs to filter IPv6 traffic:

source ipv6
destination ipv6
ipv6
nh (next header)
flow-label
source tcp port
destination tcp port
source udp port
destination udp port

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.

IPv6 ACLs are not supported on A1 NI modules. Use the **show ni** command to verify the version of the NI module. Contact your Alcatel-Lucent support representative if you are using A1 boards.

ACL & Layer 3 Security

The following additional ACL features are available for improving network security and preventing malicious activity on the network:

- **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 **icmcode**.
- **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. Note that this group is not supported on the OmniSwitch 6800.
- **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, IS-IS, 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. Note that configuring a UserPorts profile is not supported on the OmniSwitch 6800.
- **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. Note that this group is not supported on the OmniSwitch 6800.

ACL Manager

The Access Control List Manager (ACLMAN) is a function of the Quality of Service (QoS) application that provides an interactive shell for using common industry syntax to create ACLs.

Commands entered using the ACLMAN shell are interpreted and converted to Alcatel-Lucent CLI syntax that is used for creating QoS filtering policies.

This implementation of ACLMAN also provides the following features:

- Importing of text files that contain common industry ACL syntax.
- Support for both standard and extended ACLs.
- Creating ACLs on a single command line.
- The ability to assign a name, instead of a number, to an ACL or a group of ACL entries.
- Sequence numbers for named ACL statements.
- Modifying specific ACL entries without having to enter the entire ACL each time to make a change.
- The ability to add and display ACL comments.
- ACL logging extensions to display Layer 2 through 4 packet information associated with an ACL.

Account & Password Policies

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

ARP Defense Optimization

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

Detect ARP poisoning

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.

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

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 Funk/Juniper Steel Belted RADIUS server (any industry standard RADIUS server should work).
- Lightweight Directory Access Protocol (LDAP).
- Terminal Access Controller Access Control System (TACACS+).

Authentication-only servers are able to authenticate users for switch management access, but authorization (or what privileges the user has after authenticating) are determined by the switch. Authentication-only servers cannot return user privileges to the switch. The authentication-only server supported by the switch is ACE/Server, which is a part of RSA Security's SecurID product suite. RSA Security's ACE/ Agent is embedded in the switch.

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.

Authenticated VLANs

Authenticated VLANs control user access to network resources based on VLAN assignment and a user log-in process; the process is sometimes called user authentication or Layer 2 Authentication. (Another type of security is device authentication, which is set up through the use of port-binding VLAN policies or static port assignment.)

The total number of possible AVLAN users is 2K per system, not to exceed 1K per module or stackable unit. This number is a total number of users that applies to all authenticated clients, such as AVLAN and 802.1X supplicants or non-supplicants. The Omniswitch supports the use of all authentication methods and Learned Port Security (LPS) on the same port.

Layer 2 Authentication is different from Authenticated Switch Access, which is used to grant individual users access to manage the switch.

The Mac OS X 10.3.x is supported for AVLAN web authentication using JVM-v1.4.2.

Release 6.3.1 added AVLAN support for IE7 and Windows Vista. An Alcatel-Lucent certificate is required to provide this support. Please contact your customer support representative to obtain this certificate.

Automatic VLAN Containment (AVC)

In an 802.1s Multiple Spanning Tree (MST) configuration, it is possible for a port that belongs to a VLAN, which is not a member of an instance, to become the root port for that instance. This can cause a topology change that could lead to a loss of connectivity between VLANs/switches. Enabling Automatic VLAN Containment (AVC) helps to prevent this from happening by making such a port an undesirable choice for the root.

When AVC is enabled, it identifies undesirable ports and automatically configures them with an infinite path cost value.

Balancing VLANs across links according to their Multiple Spanning Tree Instance (MSTI) grouping is highly recommended to ensure that there is not a loss of connectivity during any possible topology changes. Enabling AVC on the switch is another way to prevent undesirable ports from becoming the root for an MSTI.

BGP4

The Border Gateway Protocol (BGP) is an exterior routing protocol that guarantees the loop-free exchange of routing information between autonomous systems. There are three versions of the BGP protocol— versions 2, 3, and 4. The Alcatel-Lucent implementation supports BGP version 4 as defined in RFC 1771.

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. Up to 65,000 route table entries and next hop routes can be supported by BGP.

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 new 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: 4271, 2439, 3392, 2385, 1997, 4456, 3065, 4273, 4760, and 2545.

Note that IPv6 extensions for BGP are only supported on the OmniSwitch 6850 and 9000.

Release 6.3.1 adds Webview and SNMP support.

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.

Command Line Interface (CLI)

Alcatel-Lucent's command line interface (CLI) is a text-based configuration interface that allows you to configure 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.

DHCP Relay

DHCP Relay allows you to forward DHCP broadcast requests to configurable DHCP server IP address in a routing environment.

DHCP Relay is configured using the IP helper set of commands.

Preboot Execution Environment (PXE) support was enabled by default in previous releases. Note that in this release, it is disabled by default and is now a user-configurable option using the **ip helper pxe-support** command.

DHCP Relay Agent Information Option

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

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

If the relay agent receives a DHCP packet from a client that already contains Option-82 data, the packet is dropped by default. However, it is possible to configure a DHCP Option-82 policy that directs the relay agent to drop, keep, or replace the existing Option-82 data and then forward the packet to the server.

DHCP Snooping

DHCP Snooping improves network security by filtering DHCP packets received from devices outside the network and building and maintaining a binding table (database) to log DHCP client access information. There are two levels of operation available for the DHCP Snooping feature: switch level or VLAN level.

To identify DHCP traffic that originates from outside the network, DHCP Snooping categorizes ports as either trusted or untrusted. A port is trusted if it is connected to a device inside the network, such as a DHCP server. A port is untrusted if it is connected to a device outside the network, such as a customer switch or workstation. The port trust mode is also configurable through the CLI.

Additional DHCP Snooping functionality includes the following:

- **Layer 2 DHCP Snooping**—Applies DHCP Snooping functionality to bridged DHCP client/server broadcasts without using the relay agent or requiring an IP interface on the client/server VLAN.
- **IP Source Filtering**—Restricts DHCP Snooping port traffic to only packets that contain the client source MAC address and IP address obtained from the DHCP lease information. The DHCP Snooping binding table is used to verify the client lease information for the port that is enabled for IP source filtering.

- **Rate Limiting**—Limits the number of DHCP packets on a port. This functionality is provided using the QoS application to configure ACLs for the port.
- **User-configurable Option 82 Suboption Format**—Allows the user to specify the type of information (switch base MAC address, system name, or user-defined string) that is inserted into the Circuit ID and Remote ID suboptions of the Option-82 field. This functionality only applies when DHCP Snooping Option-82 Data Insertion is enabled.

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.

Dynamic VLAN Assignment (Mobility)

Dynamic assignment applies only to mobile ports and requires the additional configuration of VLAN rules. When traffic is received on a mobile port, the packets are examined to determine if their content matches any VLAN rules configured on the switch. Rules are defined by specifying a port, MAC address, protocol, network address, binding, or DHCP criteria to capture certain types of network device traffic. It is also possible to define multiple rules for the same VLAN. A mobile port is assigned to a VLAN if its traffic matches any one VLAN rule.

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.

End User Partitioning (EUPM)

EUPM is used for customer login accounts that are configured with end-user profiles (rather than functional privileges specified by partitioned management). Profiles specify command areas as well as VLAN and/or port ranges to which the user has access. These profiles are typically used for end users rather than network administrators.

Traffic Anomaly Detection (TAD)

The Traffic Anomaly Detection (TAD) feature, also referred to as Network Security, is used to detect anomalies through statistical analysis of network traffic. It can be used to detect network attacks by observing the patterns of a port through ingress and egress packets. Anomalies occur in network traffic when the traffic patterns in a network do not meet the expectations. Such anomalies are detected in real time network traffic and can be logged, generate SNMP traps, or result in disabling the anomalous port automatically.

Network Security provides the following capabilities:

- Real time network traffic monitoring.
- Dynamic anomaly detection.

- Dynamic anomalous port quarantining.

Ethernet Interfaces

Ethernet and Gigabit Ethernet port software is responsible for a variety of functions that support Ethernet, Gigabit, and 10 Gigabit Ethernet ports. These functions include initialization of ports, notifying other software modules when a port goes down, configuration of basic line parameters, gathering of statistics for Ethernet and Gigabit Ethernet ports, and responding to administrative enable/disable requests.

Configurable parameters include: autonegotiation (copper ports 10/100/1000), trap port link messages, flood control, line speed, duplex mode, inter-frame gap, resetting statistics counters, and maximum and peak flood rates.

Flood control is configurable on ingress interfaces (flood rate and including/excluding multicast).

Ethernet OAM

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

Release 6.3.1 includes support for the IEEE 802.1ag draft 7.0 standard.

Generic 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, and TACACS) or destined for a user-defined service port can be forwarded to a maximum of 256 VLANs on the switch.

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.

GVRP

The GARP VLAN Registration Protocol (GVRP), a protocol compliant with 802.1Q, dynamically learns and further propagates VLAN membership information across a bridged network. GVRP dynamically maintains and updates the registration and de-registration of VLANs and prunes unnecessary broadcast and unicast traffic. Through propagation of GVRP information, a device is continuously able to update its knowledge of the set of VLANs that currently have active members and of the ports through which those members can be reached. With GVRP, a single switch is manually configured with all the desired VLANs for the network, and all other switches on the

network dynamically learn those VLANs. An end station can be plugged into any switch and can be connected to its desired VLAN. However, for end stations to make use of GVRP, they need Network Interface Cards (NIC) aware of GVRP.

Release 6.3.1 added the following:

- OS9000 Support
- Generate a trap if the number of dynamic VLANs exceeds the maximum threshold configured for GVRP.

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.

HTTP/HTTPS Port Configuration

The default HTTP port and the default Secure HTTP (HTTPS) port can be configured for the embedded Web server in the switch.

IP/IP Tunneling

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

IP Multicast VLAN

IP Multicast VLAN involves the creation of separate, dedicated VLANs constructed specifically for multicast traffic distribution. These distribution VLANs connect to the nearest multicast router and support multicast traffic only. The IP Multicast feature works in both the enterprise environment and the VLAN Stacking environment. The ports are separately classified as VLAN stacking ports or as legacy ports (Fixed ports/Tagged Ports). To ascertain that data flow is limited to either the VLAN Stacking domain or the enterprise domain, VLAN Stacking ports must be members of only the VLAN Stacking VLANs, while the normal legacy ports must be members of only enterprise mode VLANs.

Release 6.3.1 adds support for multiple sender ports.

Interswitch Protocol (AMAP)

Alcatel-Lucent Interswitch Protocols (AIP) are used to discover adjacent switches and retain mobile port information across switches. By default, AMAP is enabled.

Alcatel-Lucent Mapping Adjacency Protocol (AMAP) is used to discover the network topology of Alcatel-Lucent switches in a particular installation. Using this protocol, each switch determines which switches are adjacent to it by sending and responding to Hello update packets. For the purposes of AMAP, adjacent switches are those that:

- Have a Spanning Tree path between them
- Do not have any switch between them on the Spanning Tree path that has AMAP enabled

IPv4 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)
- Telnet - Client and server
- File Transfer Protocol (FTP) – Client and server
- Address Resolution Protocol (ARP)
- Internet Control Message Protocol (ICMP)
- RIP I / RIP II
- 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.

OmniSwitch 6850 and 9000 switches support hardware routing/flooding to static ARP with multicast MAC address.

The switch operates only in single MAC router mode. In this mode, each router VLAN is assigned the same MAC address, which is the base chassis MAC address for the switch.

IPv6 Support

IPv6 (documented in RFC 2460) is designed as a successor to IPv4 and is supported on the OmniSwitch 6850 and 9000. 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
- Ping, traceroute
- DNS client using Authority records
- Telnetv6 - Client and server
- File Transfer Protocol (FTPv6) – Client and server
- SSHv6 – Client and Server

OmniSwitch 6850 and 9000 switches support hardware-based IPv6 routing.

The switch operates only in single MAC router mode. In this mode, each router VLAN is assigned the same MAC address, which is the base chassis MAC address for 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 - OS6800/OS6850/OS9000
- Invalid IP Attack - OS6850/OS9000
- Multicast IP and MAC Address Mismatch - OS6850/OS9000
- Ping Overload - OS6850/OS9000
- Packets with loopback source IP address - OS6850/OS9000

IP Multicast Switching (IPMS)

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 OmniSwitch 9000 Series 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. IPMS is supported on IPv4 and IPv6 (MLD) on the OmniSwitch 6850 Series and OmniSwitch 9000 Series. The OmniSwitch 6800 Series only supports IPMS for IPv4.

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.

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.

IPX Routing

The Internet Packet Exchange (IPX) protocol, developed by Novell for NetWare, is a Layer 3 protocol used to route packets through IPX networks. (NetWare is Novell's network server operating system.) This implementation of IPX routing is software based with limited performance.

IPX specifies a connectionless datagram similar to the IP packet of TCP/IP networks. An IPX network address consists of two parts: a network number and a node number. The IPX network number is assigned by the network administrator. The node number is the Media Access Control (MAC) address for a network interface in the end node.

IS-IS

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

6.3.1 adds support for the OS-9000 as well as simple and MD5 authentication.

L2 DHCP Snooping

By default, DHCP broadcasts are flooded on the default VLAN for the client/server port. If the DHCP client and server are both members of the same VLAN domain, the broadcast packets from these sources are bridged as Layer 2 traffic and not processed by the relay agent.

The Omnswitch provides enhancements to DHCP Snooping to allow application of DHCP Snooping functionality to bridged DHCP client/server broadcasts without using the relay agent or requiring an IP interface on the client/server VLAN.

When DHCP Snooping is enabled at the switch level or for an individual VLAN, DHCP Snooping functionality is automatically applied to Layer 2 traffic. When DHCP Snooping is disabled at the switch level or disabled on the last VLAN to have snooping enabled on the switch, DHCP Snooping functionality is no longer applied to Layer 2 or Layer 3 traffic.

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. -

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

LPS has the following limitations:

- You cannot configure LPS on 10 Gigabit ports.
- You cannot configure LPS on link aggregate ports.

Link Aggregation (static & 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 32 link aggregation groups that can consist of 2, 4, or 8 Ethernetports.
- **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.

Alcatel-Lucent's link aggregation software allows you to configure the following two different types of link aggregation groups:

- Static link aggregate groups
- Dynamic (802.3ad) link aggregate groups

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).

Release 6.3.1 adds support for IPv6 via a new set of CLI commands and SNMP MIBs. Existing PIM configurations will be converted to the new CLI commands when a switch is upgraded to 6.3.1.

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. It provides client time accuracies within half a second on LANs and WANs relative to a primary server synchronized to Universal Coordinated Time (UTC) (via a Global Positioning Service receiver, for example).

OSPFv2/OSPFv3

Open Shortest Path First version 3 (OSPFv3) is available. 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.

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. Note that OSPFv3 does not support graceful restart.

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.

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.

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.

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.

Port Monitoring

The Port Monitoring feature allows you to examine packets to and from a specific Ethernet port (either ingress or egress). You can select to dump captured data to a file, which can be up to 140K. Once a file is captured, you can FTP it to a Protocol Analyzer or PC for viewing. The OmniSwitch 9000 supports one session per switch.

By default, the switch will create a data file called “pmonitor.enc” in flash memory. When the 140K limit is reached the switch will begin overwriting the data starting with the oldest captured data. However, you can configure the switch so it will not overwrite the data file. In addition, you can configure additional port monitoring files as long as you have enough room in flash memory. You cannot configure port mirroring and port monitoring on the same NI module.

Power over Ethernet (PoE)

The Power over Ethernet (PoE) software is supported on the OS6850-P24, OS6850-P24X, OS6850-P48, and OS6850-P48X stackable switches and the OS9-GNI-P24 module. PoE provides inline power directly from the switch’s Ethernet ports. From these RJ-45 ports the devices receive both electrical power and data flow. PoE detects power based on PSE devices and not on class.

PoE supports both IEEE 802.3af and non-IEEE 802.3af standards. The default inline power allotted for each port is 15400 Milliwatts. The minimum inline power allotted for a port is 3000 Milliwatts and the maximum is 16000 Milliwatts (OS6850) and 18000 Milliwatts (OS9000).

The maximum PoE power that a 510w power-supply (OS6850/OS9600) can provide is approximately 390 watts. A 360w power-supply (OS6850/OS9600) can provide approximately 240 watts of PoE power. The OS-IP-Shelf power supplies (OS9000) can provide approximately 600 watts of PoE power. The OS-IP- Shelf supports up to four power supplies, so a total of approximately 2400 watts is possible.

The redundant power supply for PoE is only for backup. If the primary power supply fails, then PoE can switch over seamlessly to the backup power supply.

PVST+ Interoperability

The current Alcatel-Lucent *IxI* 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.

Quality of Service (QoS)

Alcatel-Lucent's QoS software provides a way to manipulate flows coming through the switch based on user-configured policies. The flow manipulation (generally referred to as *Quality of Service* or *QoS*) may be as simple as allowing/denying traffic, or as complicated as remapping 802.1p bits from a Layer 2 network to ToS values in a Layer 3 network. QoS can support up to 2048 policies and it is hardware-based on the first packet. OmniSwitch 6850/9000 switches truly support 8 queues per port.

QoS is implemented on the switch through the use of policies, created on the switch or stored in PolicyView. While policies may be used in many different network scenarios, there are several typical types:

Basic QoS—includes traffic prioritization and bandwidth shaping

- **802.1p/ToS/DSCP**—includes policies for marking and mapping
- **Release 6.3.1 added support for DSCP Ranges.**
- **Policy Based Routing (PBR)**—includes policies for redirecting routed traffic
- **Access Control Lists (ACLs)**—ACLs are a specific type of QoS policy used for Layer 2, Layer 3/4, and multicast filtering.

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.

Note: When automatic NMS prioritization is enabled, QoS policies that specify priority are not applied to the NMS traffic. Other QoS policies, however, are applied to this type of traffic as usual. If a policy specifies rate limiting, then the policy with the lowest rate limiting value is applied.

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

00-80-9F-54-xx-xx to 00-80-9F-64-xx-xx

00-80-9F-66-xx-xx to 00-80-9F-6F-xx-xx.

Third-party devices can be added to this group as well.

Note: When automatic NMS prioritization is enabled, QoS policies that specify priority are not applied to the NMS traffic. Other QoS policies, however, are applied to this type of traffic as usual.

BPDU Shutdown Ports - The BPDUShutdownPorts group is a special QoS port group that identifies its members as ports that should not receive BPDUs. If a BPDU is received on one of these ports, the port is administratively disabled.

Note that the BPDUShutdownPorts group is *not* supported on the OmniSwitch 6850 Series or the OmniSwitch 9000 Series. On these switches, it is possible to configure a global UserPorts profile, as described in “ACL & Layer 3 Security”, to monitor BPDU on user ports. Such a profile also determines whether user ports will filter BPDU or will administratively shutdown when BPDU are received on the port. Note that this functionality only applies to ports that are designated as members of the UserPorts port group.

A port configured to administratively shutdown when BPDU are detected will generate an inferior BPDU every 5 seconds. This will prevent loops in the network if two BPDU shutdown ports are accidentally bridged together either through an external loop or through a hub, since both ports would be receiving inferior BPDUs.

Policy Based Mirroring

This feature enhances the current 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.

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-based 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.

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. However, on the OmniSwitch 6850 and 9000 switches, configuring minimum and maximum egress bandwidth is supported on a per COS queue basis for each port. Also note that configuring the maximum bandwidth for ingress ports is not supported on the OmniSwitch 6800.

Quarantine Manager and Remediation (QMR)

Quarantine Manager and Remediation (QMR) is a switch-based application that interacts with the OmniVista Quarantine Manager (OVQM) application to restrict the network access of quarantined clients and provide a remediation path for such clients to regain their network access. This functionality is driven by OVQM, but the following QMR components are configured through QoS CLI commands:

Quarantined MAC address group. This is a reserved QoS MAC address group that contains the MAC addresses of clients that OVQM has quarantined and that are candidates for remediation.

Remediation server and exception subnet group. This is a reserved QoS network group, called “alaExceptionSubnet”, that is configured with the IP address of a remediation server and any subnets to which a quarantined client is allowed access. The quarantined client is redirected to the remediation server to obtain updates and correct its quarantined state.

Remediation server URL. This is the URL for the remediation server. Note that this done in addition to specifying the server IP address in the “alaExceptionSubnet” network group.

Quarantined Page. When a client is quarantined and a remediation server URL is not configured, QMR can send a Quarantine Page to notify the client of its quarantined state.

HTTP proxy port group. This is a known QoS service group, called “alaHTTPProxy”, that specifies the HTTP port to which quarantined client traffic is redirected for remediation. The default HTTP port used is TCP 80 and TCP 8080.

Note that configuring QMR and QoS inner VLAN or inner 802.1p policies is mutually exclusive. QMR overlays the inner VLAN tag, thus creating a conflict with related QoS policies. This is also true with QMR and VLAN Stacking services.

QMR is activated when OVQM populates the MAC address group on the LDAP server with quarantined MAC addresses. If VLAN Stacking services or QoS inner VLAN/802.1p policies are configured on the switch, QMR will not activate.

NOTE: This feature is designed to work in conjunction with OmniVista’s Quarantine Manager application. Refer to the OmniVista documentation for a detailed overview of the Quarantine Manager application.

Within OmniVista’s Quarantine Manager application, if a MAC is added or removed from the quarantined group, or when an IP address is added or removed from the IP DA remediation, OmniVista will trigger the configured switches to perform a “recache” action. The switches will then query OmniVista’s LDAP database and “pull” the addresses from the database, these addresses will then be added or removed from the switch’s quarantined or remediation group.

Remote Port Mirroring (802.1Q based)

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

- Remote port mirroring is supported only on OS6850 and OS9000 switches.
- 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.
- BPDU mirroring will be disabled by default on all OS6850s.
- BPDU mirroring will be disabled by default on all OS9000s with B2 revision ASICs. (Contact Service and Support to enable)
- BPDU mirroring will be enabled by default on all OS9000s with A0/A1 revision ASICs.
- Source learning must be disabled or overridden on the ports belonging to the remote port mirroring VLAN on the intermediate and destination switches.
- On OS9000 and OS6850 switches the QoS redirect feature can be used to override source learning.

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 6800/6850/9000 switches support 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.

Release 6.3.1 added ECMP capability for up to 4 paths.

RIPng

The OmniSwitch 6850/9000 switches support 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.

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.

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.

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.

This feature supports basic RMON 4 group implementation in compliance with RFC 2819, including the **Ethernet Statistics**, **History** (Control & Statistics), **Alarms**, and **Events** groups.

Router Discovery Protocol (RDP)

The Router Discovery Protocol (RDP) is an extension of ICMP that allows end hosts to discover routers on their networks. The implementation of RDP supports the router requirements as defined in RFC 1256. Using RDP, hosts attached to multicast or broadcast networks send solicitation messages when they start up. Routers respond to solicitation messages with an advertisement message that contains the router IP addresses. In addition, routers send advertisement messages when their RDP interface becomes active and then subsequently at random intervals.

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. By default, local routes always have precedence.

RRSTP

Ring Rapid Spanning Tree Protocol (RRSTP) is complimentary to either the Rapid Spanning Tree (RSTP) or the Multiple Spanning Tree Protocol (MSTP) but is designed to enhance convergence time in a ring configuration on a link failure. Note that RRSTP is supported only in a ring topology where switches are connected point to point. In addition, there can be no alternate connections for the same instance between any two switches within a ring topology. RRSTP reduces convergence time by finding the bridge that hosts the alternate (ALT) port and immediately changing the ALT port state to forwarding without altering the port state. This process quickly enables the data path. The RRSTP frame travels from the point of failure to the ALT port in both directions. The MAC addresses corresponding to the ports in the ring are flushed to make the data path convergence time much faster. While RRSTP is already reacting to the loss of connectivity, the standard BPDU carrying the information about the link failure is processed in normal fashion at each hop. When this BPDU reaches the bridge whose ALT port is now in the "ALT FWD" state, due to RRSTP frame processing, it updates the state of the two ports in the ring as per the STP standard.

RRSTP is only supported when the switch is configured in Flat mode (RRSTP or MSTP).

6.3.1 adds support for the OS9000 and VLAN Stacking .

Secure Copy (SCP)

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

Secure Shell (SSH)

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.

When used as an SSH Server, the following SSH Software is supported on the indicated operating systems:

SSH Software	Supported Operating Systems
OpenSSH	Sun Solaris, Mac OSX, Linux Red Hat
F-Secure	Sun Solaris, Win 2000, Win XP
SSH-Communication	Sun Solaris, Win 2000, Win XP, Linux Red Hat
PuTTY	Win 2000, Win XP
MAC-SSH	Mac OSX

When used as an SSH Client, the following SSH Software is supported on the indicated operating systems:

SSH Software	Supported Operating Systems
OpenSSH	Sun Solaris, Linux Red Hat, AOS
F-Secure	Sun Solaris, Win 2000
SSH-Communication	Sun Solaris, Win 2000, Win XP, Linux Red Hat

Secure Shell (SSH) Public Key Authentication

DSA public key authentication is supported when using PuTTY SSH software to generate the private and public key for the client and to access the switch. It is now possible to enforce the use of public key authentication only on the switch. By default, both password and public key authentication are allowed.

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. OmniSwitch 6850/9000 switches operate at wire speed to process client requests addressed to the VIP of an SLB cluster or classified by a QoS policy condition and send them to the physical servers within the cluster.

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

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).

Smart Continuous Switching - OmniSwitch 6800/OmniSwitch 6850

In stacked configurations, one switch is designated as the primary “management module” for the stack. Because the stack can be thought of as a virtual chassis, the role of this primary management switch is to monitor and manage the functions of the entire stack.

Similar to chassis-based switches, the stack also includes a secondary, or backup, management module. A stack’s secondary switch immediately takes over management functions in the event of a primary switch failure.

All switches in the stack, besides the primary and secondary switch, are considered idle or in pass-through. Idle switches act like Network Interface (NI) modules in chassis-based switches.

The stack provides support for all idle switches during primary switch failover. In other words, if the primary switch in the stack fails or goes offline for any reason, all idle switches will continue data transmission during the secondary switch’s takeover process..

MAC Retention - The MAC Retention functionality is implemented to enhance Smart Continuous Switching for stackable products by retaining the base MAC address of the primary stack element during a takeover. As a result, both L2 and L3 traffic as well as the associated control protocols (e.g. routing protocols, spanning tree) will be minimally affected during takeover.

Release 6.3.1 added enhancements for avoiding duplicate MAC scenarios. If the primary element is not returned to the stack after a preset time, a trap will be generated indicating the possibility of a duplicate MAC. A duplicate MAC scenario would occur if the primary element was put back into the network since the stack has retained the primary element's MAC address.

Smart Continuous Switching - OmniSwitch 9000

Each OS9000 CMM module contains hardware and software elements to provide management functions for the OS9000 system. The OS9000 CMM module also contains the switch fabric for the OS9000 system. User data flowing from one NI module to another passes through the switch fabric.

The OS9700 will operate with one or two CMM modules installed. The OS9600 operates with one CMM.

If there are two CMM modules in an OS9700, one management processor is considered “primary” and is actively managing the system. The other management processor is considered “secondary” and remains ready to quickly take over management in the event of hardware or software failure on the primary. In the event of a failure, the two processors exchange roles and the secondary takes over as primary.

The switch fabric on the CMM operates independently of the management processor. If there are two CMM modules installed in an OS9700, both fabric modules are normally active. Two CMM modules must be installed in the OS9700 to provide full fabric capacity. However, note that only the one CMM module in the OS9600 provides full fabric capacity.

If there is one CMM module installed in an OS9700, then there is a single management processor, but there is no “secondary” CMM. Hardware or software failures in the CMM may result in a system reboot. The System fabric capacity on an OS9700 is one half of the fabric capacity of a dual CMM system.

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. OmniSwitch 9000 switches support SNMPv1, SNMPv2, and SNMPv3.

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.

MAC Address Mode

There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K per OmniSwitch 9000 chassis.

Note that distributed MAC address mode is not supported on the OmniSwitch 6800 Series or the OmniSwitch 6850 Series. These switches operate only in the synchronized mode.

Software Rollback

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 the working (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.

Spanning Tree

In addition to the Q2005 version of MSTP, the Alcatel-Lucent Spanning Tree implementation also provides support for the 802.1w Rapid Spanning Tree Algorithm and Protocol (RSTP) and the 802.1D Spanning Tree Algorithm and Protocol (STP). All three supported protocols ensure that there is always only one data path between any two switches for a given Spanning Tree instance to prevent network loops.

Q2005 (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.

Syslog to Multiple Hosts

Sending syslog files to multiple hosts is allowed. It is possible to specify up to a maximum of four servers.

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.

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.

User Definable Loopback Interface

Loopback0 is the name assigned to an IP interface to identify a consistent address for network management purposes. The Loopback0 interface is not bound to any VLAN, therefore it always remains operationally active. This differs from other IP interfaces, such that if there are no active ports in the VLAN, all IP interfaces associated with that VLAN are not active. In addition, the Loopback0 interface provides a unique IP address for the switch that is easily identifiable to network management applications.

User Network Profiles

This feature provides the capability to have "Roles" assigned to users during authentication. The 6.3.1 release allows for a VLAN to be associated to a role, users matching the role will automatically be assigned to that VLAN. The role should be configured to match the Filter-ID attribute being returned by the RADIUS server.

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.
- Defining VLAN IPX router interfaces to enable routing of VLAN IPX traffic.
- Enabling or disabling unique MAC address assignments for each router VLAN defined.
- Displaying VLAN configuration information.

Up to 4094 VLANs for Flat Spanning Tree mode and 252 VLANs for 1x1 Spanning Tree mode are supported. In addition, it is also possible to specify a range of VLAN IDs when creating or deleting VLANs and/or configuring VLAN parameters, such as Spanning Tree bridge values.

VLAN Stacking and Translation

VLAN Stacking 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..

VLAN Stacking Legacy and Eservice Modes - The VLAN Stacking application operates in one of two modes: Legacy and Eservice. The two modes basically differ in how VLAN Stacking is configured, with the Eservice mode offering the following additional enhancements that are not available in the Legacy mode:

- 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.
- GVRP control frame processing.
- Profiles for saving and applying traffic engineering parameter values.

Configuring VLAN Stacking in the Legacy mode consists of using a port or port-VLAN level approach to tunneling customer traffic. Configuring VLAN Stacking in the Eservices mode consists of using an approach based on defining an Ethernet service to tunnel customer traffic. Both modes are exclusive in that the switch can only operate in one mode or the other. In addition, each mode has its own unique CLI command syntax.

VRRPv2/VRRPv3

The Virtual Router Redundancy Protocol version 3 (VRRPv3) implementation is based on the latest Internet-Draft for VRRP for IPv6. VRRP version 2 (VRRPv2) is based on RFC 2338.

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. Authentication is not supported.

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.

Note that VRRPv3 is not available on the OmniSwitch 6800 Series. VRRPv2 is available on all supported OmniSwitch platforms in this release.

Global VRRP Configuration

Release 6.3.1 added the following capabilities for VRRP2 only:

- Globally enable or disable all or a range of VRRP instances.
- View or configure default values such as priority, preempt, or advertising interval on all or a group of VRRP instances.

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:

- Internet Explorer 6.0 and later for Windows NT, 2000, XP, 2003
- Firefox 2.0 for Windows and Solaris SunOS 5.10
- Windows Vista

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

Supported Traps

The following traps are supported in 6.3.1.R01:

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

			the established state.
19	bgpBackwardTransition	all	This trap is generated when the BGP router port has moved from a more active to a less active state.
20	esmDrvTrapDropsLink	all	This trap is sent when the Ethernet code drops the link because of excessive errors.
21	pimNeighborLoss	all	Signifies the loss of adjacency with a neighbor device. This trap is generated when the neighbor time expires and the switch has no other neighbors on the same interface with a lower IP address than itself.
22	dvmrpNeighborLoss	all	A 2-way adjacency relationship with a neighbor has been lost. This trap is generated when the neighbor state changes from “active” to “one-way,” “ignoring” or “down.” The trap is sent only when the switch has no other neighbors on the same interface with a lower IP address than itself.
23	dvmrpNeighborNotPruning	all	A non-pruning neighbor has been detected in an implementation-dependent manner. This trap is generated at most once per generation ID of the neighbor. For example, it should be generated at the time a neighbor is first heard from if the prune bit is not set. It should also be generated if the local system has the ability to tell that a neighbor which sets the prune bit is not pruning any branches over an extended period of time. The trap should be generated if the router has no other neighbors on the same interface with a lower IP address than itself.
24	risingAlarm	all	An Ethernet statistical variable has exceeded its rising threshold. The variable’s rising threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
25	fallingAlarm	all	An Ethernet statistical variable has dipped below its falling threshold. The variable’s falling threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
26	stpNewRoot	all	Sent by a bridge that became the new

			root of the spanning tree.
27	stpRootPortChange	all	A root port has changed for a spanning tree bridge. The root port is the port that offers the lowest cost path from this bridge to the root bridge.
28	mirrorConfigError	all	The mirroring configuration failed on an NI. This trap is sent when any NI fails to configure mirroring. Due to this error, port mirroring session will be terminated.
29	mirrorUnlikeNi	all	The mirroring configuration is deleted due to the swapping of different NI board type. The Port Mirroring session which was active on a slot cannot continue with the insertion of different NI type in the same slot.
30	slPCAMStatusTrap	all	The trap status of the Layer 2 pseudo-CAM for this NI.
31	unused	—	
32	unused	—	
33	slbTrapOperStatus	—	A change occurred in the operational status of the server load balancing entity.
34	ifMauJabberTrap	all	This trap is sent whenever a managed interface MAU enters the jabber state.
35	sessionAuthenticationTrap	all	An authentication failure trap is sent each time a user authentication is refused.
36	trapAbsorptionTrap	all	The absorption trap is sent when a trap has been absorbed at least once.
37	alaStackMgrDuplicateSlotTrap	—	Two or more slots claim to have the same slot number.
38	alaStackMgrNeighborChangeTrap	—	Indicates whether or not the stack is in loop.
39	alaStackMgrRoleChangeTrap	—	Indicates that a new primary or secondary stack is elected.
40	lpsViolationTrap	all	A Learned Port Security (LPS) violation has occurred.
41	alaDoSTrap	all	Indicates that the sending agent has received a Denial of Service (DoS) attack.
42	gmBindRuleViolation	all	Occurs whenever a binding rule which has been configured gets violated.
43	unused	—	
44	unused	—	
45	unused	—	
46	unused	—	
47	pethPsePortOnOff	—	Indicates if power inline port is or is not delivering power to the a power

			inline device.
48	pethPsePortPowerMaintenanceStatus	—	Indicates the status of the power maintenance signature for inline power.
49	pethMainPowerUsageOn	—	Indicates that the power inline usage is above the threshold.
50	pethMainPowerUsageOff	—	Indicates that the power inline usage is below the threshold.
51	ospfNbrStateChange	all	Indicates a state change of the neighbor relationship.
52	ospfVirtNbrStateChange	all	Indicates a state change of the virtual neighbor relationship.
53	httpServerDoSAttackTrap	all	This trap is sent to management station(s) when the HTTP server is under Denial of Service attack. The HTTP and HTTPS connections are sampled at a 15 second interval. This trap is sent every 1 minute while the HTTP server detects it is under attack.
54	alaStackMgrDuplicateRoleTrap	—	The element identified by alaStackMgrSlotNINumber detected the presence of two elements with the same primary or secondary role as specified by alaStackMgrChasRole on the stack.
55	alaStackMgrClearedSlotTrap	—	The element identified by alaStackMgrSlotNINumber will enter the pass through mode because its operational slot was cleared with immediate effect.
56	alaStackMgrOutOfSlotsTrap		One element of the stack will enter the pass through mode because there are no slot numbers available to be assigned to this element.
57	alaStackMgrOutOfTokensTrap		The element identified by alaStackMgrSlotNINumber will enter the pass through mode because there are no tokens available to be assigned to this element.
58	alaStackMgrOutOfPassThruSlotsTrap		There are no pass through slots available to be assigned to an element that is supposed to enter the pass through mode.
59	gmHwVlanRuleTableOverloadAlert	all	An overload trap occurs whenever a new entry to the hardware VLAN rule table gets dropped due to the overload of the table.
60	lnkaggAggUp	all	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.

61	lnkaggAggDown	all	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
62	lnkaggPortJoin	all	This trap is sent when any given port of the link aggregate group goes to the attached state.
63	lnkaggPortLeave	all	This trap is sent when any given port detaches from the link aggregate group.
64	lnkaggPortRemove	all	This trap is sent when any given port of the link aggregate group is removed due to an invalid configuration.
65	pktDrop	all	The pktDrop trap indicates that the sending agent has dropped certain packets (to blocked IP ports, from spoofed addresses, etc.).
66	monitorFileWritten	all	A File Written Trap is sent when the amount of data requested by the user has been written by the port monitoring instance.
67	alaVrrp3TrapProtoError	all	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
68	alaVrrp3TrapNewMaster	all	The SNMP agent has transferred from the backup state to the master state.
69	gmHwMixModeSubnetRuleTableOverloadAlert	all	A subnet overload trap occurs in mixed mode whenever a new entry to the HW subnet rule table gets dropped in OS6800 due to the overload of the table.
70	pethPwrSupplyConflict	all	Power supply type conflict trap.
71	pethPwrSupplyNotSupported	all	Power supply not supported trap.
72	chassisTrapsPossibleDuplicateMac	6800/6850	The old PRIMARY element cannot be detected in the stack. There is a possibility of a duplicate MAC address in the network
73	vRtrIisisDatabaseOverload	all	This notification is generated when the system enters or leaves the Overload state.
74	vRtrIisisManualAddressDrops	all	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
75	vRtrIisisCorruptedLSPDetected	all	This notification is generated when an LSP that was stored in memory has become corrupted.
76	vRtrIisisMaxSeqExceedAttempt	all	Generated when the sequence number on an LSP wraps the 32 bit sequence counter

77	vRtrIsisIDLenMismatch	all	Need Desc. A notification sent when a PDU is received with a different value of the System ID Length.
78	vRtrIsisMaxAreaAddrsMismatch	all	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
79	vRtrIsisOwnLSPPurge	all	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
80	vRtrIsisSequenceNumberSkip	all	When we receive an LSP is received without a System ID and different contents.
81	vRtrIsisAutTypeFail	all	A notification sent when a PDU is received with the wrong authentication type field.
82	vRtrIsisAuthFail	all	A notification sent when a PDU is received with an incorrect authentication information field.
83	vRtrIsisVersionSkew	all	A notification sent when a Hello PDU is received from an IS running a different version of the protocol.
84	vRtrIsisAreaMismatch	all	A notification sent when a Hello PDU is received from an IS which does not share any area address.
85	vRtrIsisRejectedAdjacency	all	A notification sent when a Hello PDU is received from an IS, but does not establish an adjacency due to a lack of resources.
86	vRtrIsisLSPTooLargeToPropagate	all	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
87	vRtrIsisOrigLSPBufSizeMismatch	all	A notification sent when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for the originating L1LSPBufferSize or originating L2LSPBufferSize respectively. Also when a Level 1 LSP or Level 2 LSP is received containing the originating LSPBufferSize option and the value in the PDU option field does not match the local value for originatingL1LSPBufferSize or originatingL2LSPBufferSize respectively.
88	vRtrIsisProtoSuppMismatch	all	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
89	vRtrIsisAdjacencyChange	all	A notification sent when an adjacency changes state, entering or leaving state

			up. The first 6 bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
90	vRtrIsisCircIdExhausted	all	A notification sent when ISIS cannot be started on a LAN interface because a unique circId could not be assigned due to the exhaustion of the circId space.
91	vRtrIsisAdjRestartStatusChange	all	A notification sent when an adjacency's graceful restart status changes.
92	dot1agCfmFaultAlarm	all	A MEP has lost contact with one or more MEPs. A notification (fault alarm) is sent to the management entity with the OID of the MEP that has detected the fault.
93	Unused	all	-
94	lldpRemTablesChange	all	A lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes.
95	lpsPortUpAfterLearningWindowExpiredT	all	When an LPS port joins or is enabled after the Learning Window is expired, the MAC address learning on the port will be disabled, and this trap is generated as a notification.
96	alaPimNeighborLoss	all	A alaPimNeighborLoss notification signifies the loss of an adjacency with a neighbor.
97	alaPimInvalidRegister	all	An alaPimInvalidRegister notification signifies that an invalid PIM Register message was received by this device
98	alaPimInvalidJoinPrune	all	A alaPimInvalidJoinPrune notification signifies that an invalid PIM Join/Prune message was received by this device.
99	alaPimRPMappingChange	all	An alaPimRPMappingChange notification signifies a change to the active RP mapping on this device.
100	alaPimInterfaceElection	all	An alaPimInterfaceElection notification signifies that a new DR or DR has been elected on a network.
101	lpsLearnTrap	all	Generated when an LPS port learns a bridged MAC.
102	gvrpVlanLimitReachedEvent	all	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
103	alaNetSecPortTrapAnomaly	all	Trap for an anomaly detected on a port.
104	alaNetSecPortTrapQuarantine	all	Trap for an anomalous port quarantine.

105	udldStateChange	all	Generated when the state of the UDLD protocol changes.
-----	-----------------	-----	--

Unsupported Software Features

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

Feature	Platform	Software Package
OSPF Database Overflow (RFC 1765)	all	base
Flow Control 802.3x	all	base

Unsupported CLI Commands

The following CLI commands are not supported in this release of the software:

Software Feature	Unsupported CLI Commands
BGP	ip bgp redist-filter status ip bgp redist-filter ip bgp redist-filter community ip bgp redist-filter local-preference ip bgp redist-filter metric ip bgp redist-filter effect ip bgp redist-filter subnets
Chassis Mac Server	mac-range local mac-range duplicate-eprom mac-range allocate-local-only show mac-range status
Chassis Supervision	show fabric
Command Line Interface (CLI)	10 gig slot [slot] phy-a phy-b
DHCP Relay	ip helper traffic-suppression ip helper dhcp-snooping port traffic-suppression
Ethernet Interfaces	interfaces long interfaces runt interfaces runtsize
Flow Control	flow flow wait time interfaces flow
Hot Swap	reload ni [slot] # [no] power ni all
NTP	no ntp server all
OSPF	ip ospf redist status ip ospf redist ip ospf redist metric ip ospf redist metric-type ip ospf redist-filter

	<p>ip ospf redist-filter effect ip ospf redist-filter metric ip ospf redist-filter route-tag ip ospf redist-filter redist-control</p>
PIM	<p>ip pim cbsr-masklength ip pim static-rp status ip pim rp-candidate ip pim crp-address ip pim crp-expirytime ip pim crp-holdtime ip pim crp-interval ip pim crp-priority ip pim data-timeout ip pim joinprune-interval ip pim source-lifetime ip pim interface mode ip pim interface cbsr-prefernce ip pim interface max-graft-retries ip pim interface sr-ttl-threshold show ip pim rp-candidate show ip pim rp-set show ip pim nexthop show ip pim mroute</p>
QoS	<p>qos classify fragments qos flow timeout show policy classify destination interface type show policy classify source interface type</p>
RIP	<p>ip rip redist status ip rip redist ip rip redist metric ip rip redist-filter ip rip redist-filter effect ip rip redist-filter metric ip rip redist-filter route-tag ip rip redist-filter redist-control</p>
SYSTEM	<p>install</p>
VLANs	<p>vlan router mac multiple enable disable vlan binding mac-port-protocol vlan binding mac-ip vlan binding ip-port</p>

Unsupported MIBs

The following MIBs are not supported in this release of the software

Feature	MIB
Quality of Service (QoS)	IETF_P_BRIDGE
Flow Control	AlcatelIND1Port

:

Unsupported MIB Variables

MIB Name	Unsupported MIB variables
AlcatelIND1AAA	aaauProfile
AlcatelIND1Bgp	alaBgpGlobal alaBgpPeerTable alaBgpAggrTable alaBgpNetworkTable alaBgpRedistRouteTable alaBgpRouteTable alaBgpPathTable alaBgpDampTable alaBgpRouteMapTable alaBgpAspathMatchListTable alaBgpAspathPriMatchListTable alaBgpPrefixMatchListTable alaBgpCommunityMatchListTable alaBgpCommunityPriMatchListTable alaBgpDebugTable
AlcatelIND1Dot1Q	qPortVlanForceTagInternal
AlcatelIND1GroupMobility	vPortIpBRuleTable vMacIpBRuleTable vMacPortProtoBRuleTable vCustomRuleTable
AlcatelIND1Health	healthDeviceTemperatureCmmCpuLatest healthDeviceTemperatureCmmCpu1MinAvg healthDeviceTemperatureCmmCpu1HrAvg healthDeviceTemperatureCmmCpu1HrMax
AlcatelIND1Ipms	alaIpmsForwardSrcIpAddr alaIpmsForwardSrcIfIndex
AlcatelIND1LAG	alclnkaggAggEniActivate alclnkaggSlotTable
AlcatelIND1Pcam	alcatelIND1PCAMMIBObjects alaCoroL3HrePerModeTable alaCoroL3HrePerCoronadoStats Table alaCoroL3HreChangeTable

<p>AlcatelIND1Port</p>	<p>esmPortCfgLongEnable esmPortCfgRuntEnable esmPortCfgRuntSize esmPortPauseSlotTime esmPortCfgFlow</p>	<p>alcether10GigTable</p>
<p>AlcatelIND1QoS</p>	<p>alaQoSPortPdiTable alaQoSslotPcamTable alaQoSPortProtocolTable alaQoSslotProtocolTable alaQoSslotDscpTable alaQoSRuleReflexive alaQoSAppliedRuleReflexive alaQoSActionSourceRewriteIpAddr alaQoSActionSourceRewriteIpAddrStatus alaQoSActionSourceRewriteIpMask alaQoSActionTable alaQoSActionSourceRewriteNetworkGroup alaQoSActionTable alaQoSActionSourceRewriteNetworkGroupStatus alaQoSActionTable alaQoSActionDestinationRewriteIpAddr alaQoSActionTable alaQoSActionDestinationRewriteIpAddrStatus alaQoSActionTable alaQoSActionDestinationRewriteIpMask alaQoSActionTable alaQoSActionDestinationRewriteNetworkGroup alaQoSActionTable alaQoSActionDestinationRewriteNetworkGroupStatus alaQoSActionTable alaQoSActionLoadBalanceGroup alaQoSActionTable alaQoSActionLoadBalanceGroupStatus alaQoSActionTable alaQoSActionPermanentGatewayIpAddr alaQoSActionTable alaQoSActionPermanentGatewayIpAddrStatus alaQoSActionTable alaQoSActionAlternateGatewayIpAddr alaQoSActionAlternateGatewayIpAddrStatus alaQoSAppliedActionSourceRewriteIpAddr alaQoSAppliedActionSourceRewriteIpAddrStatus alaQoSAppliedActionSourceRewriteIpMask alaQoSAppliedActionSourceRewriteNetworkGroup alaQoSAppliedActionSourceRewriteNetworkGroupStatus alaQoSAppliedActionDestinationRewriteIpAddr alaQoSAppliedActionDestinationRewriteIpAddrStatus alaQoSAppliedActionDestinationRewriteIpMask alaQoSAppliedActionDestinationRewriteNetworkGroup alaQoSAppliedActionDestinationRewriteNetworkGroupStatus alaQoSAppliedActionLoadBalanceGroup alaQoSAppliedActionLoadBalanceGroupStatus alaQoSAppliedActionPermanentGatewayIpAddr alaQoSAppliedActionPermanentGatewayIpAddrStatus alaQoSAppliedActionAlternateGatewayIpAddr alaQoSAppliedActionAlternateGatewayIpAddrStatus alaQoSPortDefaultQueues alaQoSPortAppliedDefaultQueues alaQoSConfigNatTimeout alaQoSConfigAppliedNatTimeout alaQoSConfigReflexiveTimeout alaQoSConfigAppliedReflexiveTimeout</p>	

	alaQoSConfigFragmentTimeout alaQoSConfigAppliedFragmentTimeout alaQoSConfigClassifyFragments alaQoSConfigAppliedClassifyFragments	
AlcatelIND1Slb	slbFeature slbClusterTable slbServerTableg	
AlcatelIND1StackManager	alaStackMgrStatsTable	
AlcatelIND1SystemService	systemUpdateStatusTable	
AlcatelIND1VlanManager	vlanIpxNet vlanIpxEncap vlanIpxRipSapMode vlanIpxDelayTicks vlanSetMultiRtrMacStatus	vlanIpxStatus vlanSetIpxRouterCount
AlcatelIND1WebMgt	alaIND1WebMgtRFSCfgTable alaIND1WebMgtHttpPort alaIND1WebMgtHttpsPort	
IEEE_802_1X	dot1xAuthDiagTable dot1xAuthSessionStatsTable dot1xSuppConfigTable dot1xSuppStatsTable	
IETF_BGP4	bgpRcvdPathAttrTable bgp bgpPeerTable bgp4PathAttrTable	
IETF_BRIDGE	dot1dTpPortTable dot1dStaticTable	
IETF_ENTITY	entLogicalTable entLPMappingTable entAliasMappingTable	
IETF_ETHERLIKE	dot3CollTable dot3StatsSQETestErrors dot3StatsInternalMacTransmitErrors dot3StatsCarrierSenseErrors dot3StatsInternalMacReceiveErrors dot3StatsEtherChipSet dot3StatsSymbolErrors dot3ControlInUnknownOpCodes	
IETF_IF	ifRcvAddressTable ifTestTable	
IETF_IP_FORWARD_MIB	ipForwardTable	
IETF_IPMROUTE_STD	ipMrouteScopeNameTable	
IETF_MAU (RFC 2668)	rpMauTable rpJackTable broadMauBasicTable ifMauFalseCarriers ifMauTypeList ifMauAutoNegCapability ifMauAutoNegCapAdvertised	

	ifMauAutoNegCapReceived	
IETF_OSPF (RFC 1850)	ospfAreaRangeTable	
IETF_OSPF_TRAP	ospfTrapControl	
IETF-PIM	pimRPTable	
IETF_P_BRIDGE	dot1dExtBase dot1dPortCapabilitiesTable dot1dPortPriorityTable dot1dUserPriorityRegenTable dot1dTraficClassTable dot1dPortOutboundAccessPriorityTable dot1dPortGarpTable dot1dPortGmrpTable dot1dTpHCPortTable dot1dTpPortOverflowTable	
IETF_Q_BRIDGE (RFC 2674)	dot1qTpGroupTable dot1qForwardAllTable dot1qForwardUnregisteredTable dot1qStaticMulticastTable dot1qPortVlanStatisticsTable dot1qPortVlanHCStatisticsTable dot1qLearningConstraintsTable	
IETF_RIPv2	rip2IfConfDomain	
IETF_RMON	hostControlTable hostTable hostTimeTable hostTopNControlTable hostTopNTable matrixControlTable matrixSDTable matrixDSTable filterTable channelTable bufferControlTable captureBufferTable	
IETF_RS_232 (RFC 1659)	all synchronous and sdhc objects and tables rs232SyncPortTable	
IETF_SNMPv2	sysORTable snmpTrap sysORLastChange	
IETF_SNMP_COMMUNITY (RFC 2576)	snmpTargetAddrExtTable	
IETF_SNMP_NOTIFICATION (RFC 2576)	snmpNotifyTable snmpNotifyFilterProfileTable snmpNotifyFilterTable	
IETF_SNMP_PROXY (RFC 2573)	snmpProxyTable	
IETF_SNMP_TARGET (RFC 2573)	snmpTargetAddrTable snmpTargetParamsTable snmpTargetSpinLock	
IETF_SNMP_USER_BASED_SM	usmUser	

(RFC 2574)	
IETF_SNMP_VIEW_BASED_ACM (RFC 2575)	vasmMIBViews

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.

SWITCH MANAGEMENT

General

PR	Description	Workaround
119625	User 'admin' cannot change its expired password via WebView but other users work fine.	Change the password using the CLI.

CLI

PR	Description	Workaround
109841	If filtering is used in command "show ip route", only one gateway will be display if there are multiple lines display in case for ecmp routes.	There is no known workaround at this time
111704	Some debug output from "debug systrace appid aaa level debug1" can cause the telnet session to have garbage output.	Resetting the terminal will recover from this problem.
117588	CLI will provide help for unsupported parameters when using '?'.	Refer to the CLI Reference guide for supported parameters.
119091	CLI: Command "ip bgp network community" allows community string 'no-adv' that needs to be removed. It is short for 'no-advertise' and will be removed in a future release.	There is no known workaround at this time.
119288	Under very stressful conditions excessive Telnet or SSH sessions may cause management accessibility issues.	Configure the appropriate ACLs to prevent excessive Telnet or SSH connections.
119571	A delay for the vlan range command may be seen whenever a character is waiting on the socket.	Do not type ahead while waiting for VLAN range commands to complete. This should take only several seconds

Health Monitoring

PR	Description	Workaround
110452	The command "interface no l2 statistics" will cause the health monitor not updating the rtx count	After issuing "interface no l2 statistics", enter "health statistics reset". This will clear reset all health statistics to zero. They will begin to update soon thereafter.

Web Management

Feature Exceptions

WebView uses signed applets for the automatic IP reconfiguration. Those applets are signed using VeriSign Certificates that expire every year. The certificate used for Internet Explorer and Netscape

expires every August. WebView users have to validate a warning indicating that the certificate used by the applet has expired.

PR	Description	Workaround
117808	WebView Physical > Adjacencies home page drawing does not show 802.1AB LLDP devices, only AMAP.	Information available in tables under LLDP.
114491	In some WebView tables, when the “S” button for changing the sorting order is pressed and the first sorting column is selected the sorting graphic indicator number may say ‘2’ instead of ‘1’ and will continue up until ‘5’ instead of ‘4’. The functionality is still correct; however, the numbering is misleading because it skips ‘1’.	There is no known workaround at this time
113285	In WebView, two or more Internet Explorer browser windows are opened for different routers any Add, Modify, or Help windows remain opened for a router (say router A), another Add, Modify, or Help window from another router (say router B) might log out the previous router (in this case, router B’s newly opened window will log out router A).	Use only one browser window to configure a router at a time.
117411	Network security anomaly traffic pages take almost 15 to 20 seconds to complete.	CLI command “show netsec port traffic” may be used for faster access to port anomaly traffic statistics.
117429	Webview doesn’t have the capability to display per port anomaly statistics like the CLI.	CLI commands “show netsec port traffic”, “show netsec port summary”, and “show netsec port summary” commands may be used for per port anomaly statistics.
113566	In WebView Networking > UDP Relay > Services > Destination, if a service does not have a specific destination configured the service is not shown even though CLI command “show ip udp relay destination” will list it (there are no functional differences).	Refer to either Services > Configuration or Statistics (BootP / Generic Services) to see all current services.
118563	Webview networking-IP interface page is unable to add an address to Loopback0.	CLI command “ip interface” may be used instead.
116535	In WebView, whenever choosing 65 or more selections (rows in tables or rows in multiple-select drop-downs) to make an action go into effect (such as Enable, Disable, and the like in table pages or Add/Modify in corresponding pages) – only the first 64+ operations are done. Whenever using a select-all checkbox in pages that list greater than 65 rows, WebView doesn’t differentiate between the ones the operation has been done already on versus the ones it hasn’t.	Selections have to be done manually after the first 65 rows whenever using the select-all checkbox in table pages.

PR	Description	Workaround
116535	In WebView, whenever choosing 65 or more selections (rows in tables or rows in multiple-select drop-downs) to make an action go into effect (such as Enable, Disable, and the like in table pages or Add/Modify in corresponding pages) – only the first 64+ operations are done. Whenever using a select-all checkbox in pages that list greater than 65 rows, WebView doesn't differentiate between the ones the operation has been done already on versus the ones it hasn't.	Selections have to be done manually after the first 65 rows whenever using the select-all checkbox in table pages.
117581	WebView System > System Mgmt > Install "File Selection" browse and select files locks IE7 on Windows Vista.	Use Firefox on Windows Vista.
116091	WebView Layer 2 > VLAN Mgmt > VLAN Configuration > Ports > Port Association "Move ports" button fails to move more than 128 ports into a VLAN.	Move 128 ports at a time.
116171	When using Windows Vista and Microsoft Internet Explorer 7, the WebView Telnet link from the toolbar in the top right fails to connect to the device due to the telnet feature control being disabled by default (see http://msdn2.microsoft.com/en-us/library/ms537169.aspx#FEATURE_DISABLE_TELNET_PROTOCOL).	Refer to Microsoft documentation for enabling Telnet.

LAYER 2

Ethernet

PR	Description	Workaround
118967	Due to excessive system message load on takeover a number of system failures can occur resulting in one of the stack elements failing takeover. The NI will reboot if this occurs. Types of failures seen which are symptomatic of this are:task taNiEsmDrv" is not ready, taStp task failure	Enable MAC retention is recommended for systems which experience this problem. Command: MAC-RETENTION STATUS ENABLE
105080	If Gig copper SFP is plugged in a combo port then those combo ports can be used in either preferred-copper or forced mode only. Preferred-fiber mode is not supported. In forced mode, Copper combo port can be used by setting the mode to forced-copper, and to use the Gig copper SFP the mode should be set to forced-fiber.	There is no known workaround at this time

PR	Description	Workaround
105168	The model and device name are not displayed for SFPs and XFPs when the 'show ni' command is entered.	There is no known workaround at this time
106811	Show interface slot/port "SFP/XFP" field output for a port having SFP plugged in cannot differentiate between 100Fx and Bidirectional SFP & between Gig and CWDM SFP.	There is no known workaround at this time
118512	Modifying frame size from 1553 bytes to 9K bytes on OS6850 combo port (configured at speed 100 Mbps) applies only to the preferred media.	There is no known workaround at this time.
119085	On combo ports set for Preferred Copper mode, the CLI does not allow user to change fiber media speed to 100 and displays an invalid error message "Fiber GigaEthernet accepts only 1000".	Change the combo port mode from Preferred Copper to Preferred Fiber, set the fiber media speed to 100, and then change the mode back to Preferred Copper.
119253	In some cases, when layer 2 broadcasts are received from both backplane and local port on a 10G module at the same time, some of the broadcast packets might get dropped if the aggregate rate is more than 1.4Gbps.	There is no known workaround at this time.

Group Mobility

PR	Description	Workaround
98417	When a MAC is learnt as "Filtered" for one port due to a Group Mobility rule violation, and if the MAC reappears on another port, it will not be updated. That is the MAC will not be shown as filtered for the new port, but will continue to show filtered in the old port.	There is no known workaround at this time
119230	SNMP get on the port protocol ip-e2 based binding rule returns error.	Use CLI based show command.
119441	When there are multiple ether-type rules matching from the same port those rules will be matched only if the port doesn't receive any other ether-type traffic. For example, if a DECENT ether-type rule has been configured on a port and IP packets are received then DECNET would be classified in the default VLAN.	Add a separate IP subnet rule.

GVRP

PR	Description	Workaround
117983	The "configuration apply" command does not apply the GVRP configuration for link aggregate ports.	There is no known workaround at this time.
119618	Some ports may get stuck in blocking state after	There is no known workaround at this

to toggling GVRP and changing max VLAN limit. time.

Link Aggregation

PR	Description	Workaround
113407	When the Source MAC is unknown on GM ports and the Destination MAC is known, the initial few packets may be dropped if the rate of ingress traffic is high.	There is no known workaround at this time

Port Mirroring/Monitoring

PR	Description	Workaround
96691	OS6850 and OS9700 contain a default VLAN tag on all egress mirrored packets when sending only unicast untagged packets. This issue does not exist on OS6800.	
96830	Even if the mirrored port is prevented to send out flooded traffic due to spanning tree state of that particular VLAN, the mirroring port will get the flooded packet.	There is no known workaround at this time
114259	In the overwrite off mode, port monitoring will not put any frames in the pmonitor.enc file. This issue does not occur in the overwrite on mode.	There is no known workaround at this time
114260	Port Monitor fails to capture in Broadcast and Multicast packets.	There is no known workaround at this time.
114261	Port Monitor capture file contains switch's own BPDUs if Port Mobility is on.	There is no known workaround at this time
117577	When untagged traffic is egress mirrored, the mirroring port will drop some captured traffic, due to bandwidth limitations since egress mirrored traffic always gets tagged.	There is no known workaround at this time.
117590	When both policy based mirroring criteria and port based outport mirroring criteria matches for the packet, both mirroring sessions will create a unique copy to the mirroring port resulting in the mirroring port receiving two copies of the same packet.	There is no known workaround at this time.
119614	On an OS6800, egress mirrored packets generated by the hardware will have VLAN tag inserted , and egress mirrored control packets from the CPU may not have the VLAN tag, as the CPU will strip the VLAN tag. when received at the mirroring port.	There is no known workaround at this time.

Source Learning

PR	Description	Workaround
----	-------------	------------

105399	Under the following conditions, traffic is flooded instead of unicast bridging: When the chassis is operating under the distributed mode of source learning, and the traffic from slots egressing linkaggregation on different slots and the ingress and egress traffic are asymmetric.	Configure the source mac as static or use linkaggregation on the same slot.
113559	If we are adding authenticated users to the system at the rate which is less than MAC aging time (for example, endless stress test), then all MACs in the system will not be aged out until the adding is stop.	There is no workaround at this time.
118596	MAC address sometimes displays on multiple VLANs after authentication.	There is no known workaround at this time. This is only a display issue.

Spanning Tree

PR	Description	Workaround
100761	Disabling a linkAggr port while traffic is running might leave some MAC addresses learned on the primary member port of the linkAggr.	Stop the traffic first or link down/up the primary port after disabling the linkAggr.
108800	On Root Bridge down RRSTP convergence time is not sub 50 msec.	There is no known workaround at this time
119635	RRSTP takes more than 50 msec to converge after link down when RRSTP ring ports are LACP ports (fuji OS9-GNI-U24 card)	There is no known workaround at this time
118112	The current range for dot1dStpProtocolSpecification is from 1 to 3, which does not support 1w (4) or 1s (5) protocols. If the current switch is running either 1w or 1s protocol, the value of 4 or 5 will be returned for the dot1dStpProtocolSpecification object.	There is no known workaround at this time

VLAN Stacking

PR	Description	Workaround
119485	The system cannot cleanly unbind a linkagg from a SAP which does not have any CVLAN mappings. Error messages are displayed on the console.	If an SAP has a linkagg and CVLANs bound to it, and you want to unbind the linkaggs, it should be done before you unbind the CVLANs from the SAP.
119486	In Eservice mode, two CVLANs can be mapped to the same service in translation mode when the CVLANs are part of two different SAPs. The translation rule for the first configured CVLAN will be overwritten.	There is no known workaround at this time.
119575	In Eservice mode on a linkagg NNI, legacy STP/GVRP BPDUs will not be processed. The configuration to enable the processing of legacy STP/GVRP BPDUs on a linkagg NNI does not work at this time.	There is no known workaround at this time.

UDLD

PR	Description	Workaround
119593	An error is displayed when the following MIB objects are read - UdldGlobalConfigUdldProbeIntervalTimer, UdldGlobalConfigUdldDetectionPeriodTimer, UdldGlobalClearStats, UdldGlobalConfigUdldMode are read. This does not affect UDLD functionality as these are only significant for WRITE purposes.	There is no known workaround at this time.

LAYER 3

ARP

PR	Description	Workaround
116759	The statistics of ARP-Poison attacks shown by the CLI command: “show ip dos arp-poison” are likely to be inaccurate after a takeover.	There is no known workaround at this time.
119469	Static ARP with Multicast MAC doesn't work when the ingress traffic is on GNI-C48T modules. It works OK with GNI-C24 GNI-U24 and C20L modules	There is no known workaround at this time.
119573	ARP Defense will not enable when you toggle the “icmp unreachable host-unreachable” command.	Whenever configuring this parameter the switch needs to be rebooted to take affect.

IPv6

PR	Description	Workaround
103104	On OS9000 and OS6850 the IPv6 next headers are not parsed beyond the first one.	There is no known workaround at this time.
113700	The commands “show ipv6 routes” and “show ipv6 router database” will not show the exact same entries in their respective tables. When it comes to 6to4 tunnels, the latter command displays only one route entry (2002/16) associated with the tunnel. It also doesn't display any loopback addresses.	There is no known workaround at this time
119464	It is possible that a ICMP parameter problem may not be sent for unknown Destination Options present in a multicast packet even though the option code requires the transmission of the error message.	There is no known workaround at this time.
119506	IPv6 OSPFv3 cannot be configured by LDAP user with full read/write access.	Configure OSPFv3 with local authenticated user.

PIM-SM

PR	Description	Workaround
119403	Configuration of a PIM Candidate-RP is only supported on the primary IP interface, but the cli isn't currently restricting configuration on a secondary interface.	Only use the primary IP interface when configuring a PIM Candidate-RP.
119471	A switch with ipms globally enabled and which is the Designated Router on a pim interface with local IGMP SSM members may not receive non-local multicast traffic after disabling and re-enabling pim if the multicast flow still exists in the ipms tables.	Disable/re-enable pim on the incoming interface of the multicast flow.
119699	When upgrading from previous releases to 6.3.1.R01, pim adds 224.0.0.0/4 rp entry in the boot.cfg.	Remove the entry after bootup using the cli command 'no ip pim candidate-rp' or manually remove the entry from the boot.cfg.

UDP/TCP

PR	Description	Workaround
97666	TCP doesn't send a RST after receiving an unacceptable ACK in SYN-RCVD state.	There is no known workaround at this time.

UDP Relay

PR	Description	Workaround
106478	<p>When using BootP, if a write memory is done after a dynamic IP address is assigned to VLAN 1 that IP address will become permanent and is stored as a static route.</p> <p>If a takeover occurs before a write memory, the IP address is not assigned on the new primary CMM.</p>	Assign a static IP address to VLAN 1.

VRRP

PR	Description	Workaround
119633	If the VRRP advertisement granularity is modified via the "debug set vrrpAdverGranularity xx" command in the AlcatelDebug.cfg file, a CLI "write memory" command may result in a statement being generated in the boot.cfg file producing a VRRP global default interval, e.g. VRRP INTERVAL 100. This interval could then be inappropriately applied following a reboot.	Set the VRRP interval back to the appropriate value before rebooting.

Quality of Service

General

PR	Description	Workaround
94125	The OS9000 and OS6850 hardware does not support QoS or ACL rules containing destination port, destination VLAN, or destination MAC on traffic that is routed by the OS9000 or OS6850.	There is no known workaround at this time.

Server Load Balancing

PR	Description	Workaround
119449	On an OS6850, drops traffic to one of the servers after takeover	Manually flush ARP table after takeover to recover flow.

Security

802.1x

PR	Description	Workaround
118751	Device classification allow users to configure classification policy based on the 802.1x authentication status. If the supplicant failed to authenticate with the authentication server, AOS will locally authenticate the supplicant and classify the supplicant according to the failed policy to the appropriate vlan. When the supplicant is locally authenticated by the AOS, the supplicant will received a success notification from AOS. This will not work if user is using TTLS/PEAP as the authentication method. Supplicant will not be locally authenticated.	There is no known workaround at this time

AAA Services

PR	Description	Workaround
119638	A host user getting authenticated by a TACACS+ server for Debug PM family cannot do a 'show drclog' since CLIParser presents only read privileges to that CLI command.	There is no known workaround at this time.
119653	aaaServerTable MIB is not correctly populated by AAA module.	There is no known workaround at this time.

AVLAN

PR	Description	Workaround
114976	AVLAN client PC authenticated via web browser may have issue if proxy server is defined/used on client browser.	The "aaa avlan dns name" should be always defined as a singleton domain name (i.e. authent instead of authent.com), so the HTTPS

		connection is always enforced. In addition, in the web browser's setting, please make sure the "Bypass proxy server for local addresses" in the internet options's LAN setting.
114844	The Avlan pre-authentication temporary IP address will appear on the ARP table.	None. It does not affect functionality.
114844	The Avlan pre-authentication temporary IP address will appear on the ARP table.	None. It does not affect functionality.
117260	When using Windows Vista for Avlan web authentication, user might run into a problem where after authentication, the PC is not able to switch to the IP subnet of the newly classified vlan.	Refer to the following Microsoft Windows Vista tip: http://windowshelp.microsoft.com/windows/en-us/help/f941cb45-b2cd-4b39-ab87-cb9ea959f44e1033.mspx

Device Classification

PR	Description	Workaround
112338	It is possible that the supplicant is not able to transit out of the ABORTING state.	Reboot the switch.

Traffic Anomaly Detection (Network Security)

PR	Description	Workaround
119523	Network security area cannot be configured by LDAP user with full read/write access.	User local user authentication.

LDAP

PR	Description	Workaround
63005	LDAP referral service not supported	There is no known workaround at this time.

System

General

PR	Description	Workaround
93114	If an ingress packet is oversized, it is not for broadcast or multicast anymore. As a result, ingress oversized layer 2 multicast packets are counted as unicast packets.	There is no known workaround at this time
94269	The OmniSwitch counts all error packets as unicast packets in the packets received and error counters regardless of whether the packet is a unicast packet or a multicast packet. An oversize packet is defined as a packet longer than 9216 bytes. This causes the following behavior in the switch. 1) Received packets longer than 9216 bytes are counted as unicast packets AND as	There is no known workaround at this time

	error packets even if the packet is a broadcast or multicast packets. 2) For ports operating at speeds of 10 Mbps or 100 Mbps, a packet is not counted as an error packet unless it is longer than 9216 bytes.	
119258	On an OS6800, health statistics may report '0' for current CPU utilization on non-primary stacks	There is no known workaround at this time.
94866	Invalid Packets with SA or DA set to all 0's continue to get bridged by system.	There is no known workaround at this time.

Chassis Supervision

PR	Description	Workaround
114038	Time synchronization between Primary and other CMM's is offset by time required on non-primary CMM's for writing date, DST, and timezone info into eeprom.	Implement NTP client to a synchronized time source.
119609	The 'show ni' command has missing information for the 'Default and Backup' miniboot version details. Normally this version information is the same as the 'Bootrom' version, as they are normally updated together. The 'Bootrom' information is available.	There is no known workaround at this time.

NI System

PR	Description	Workaround
105646	entPhysicalModelName MIB variable returns vendor name of SFP instead of model name for an SNMP get/getNext call to this object.	There is no know workaround at this time.
119038	Adding multiple stacking modules at the same time to an existing operational stack might not bring up the new modules and integrate them as part of the existing stack.	When adding multiple stacking elements to an existing operational stack, it is advisable to add one at a time and make sure the added module is up and integrated to the operational stack before adding the next one. The other alternative is to power cycle the whole stack after connecting the stacking ports.

Stack Manager

PR	Description	Workaround
86853	Passthru mode is not supported on 6800/6850.	There is no known workaround at this time

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
Other International	818-878-4507

Email: support@ind.alcatel.com

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

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

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

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

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

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