Release Notes

OmniSwitch 9000E

Release 6.4.1.S01

These release notes accompany release 6.4.1.R01 software for the OmniSwitch 9000E (OmniSwitch 9700E/9800E) 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.

Note: The 'S01' software release is required for support of the ISSU feature. Additional details are provided in the New Software Features section.

Contents

Related Documentation

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

User manuals can be downloaded at:
http://www1.alcatel-lucent.com/enterprise/en/resource_library/ user_manuals.html.

- ***OmniSwitch 9000E Series Getting Started Guide***
Describes the hardware and software procedures for getting an OmniSwitch 9000E Series switch up and running.
- ***OmniSwitch 9000E Series Hardware User Guide***
Complete technical specifications and procedures for all OmniSwitch 9000E Series chassis, power supplies, and fans.
- ***OmniSwitch CLI Reference Guide***
Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.
- ***OmniSwitch AOS Release 6 Network Configuration Guide***
Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethernet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.
- ***OmniSwitch AOS Release 6 Switch Management Guide***
Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).
- ***OmniSwitch AOS Release 6 Advanced Routing Configuration Guide***
Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, and OSPFv3.
- ***OmniSwitch 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.

OmniSwitch 9000 and 9000E Overview

The table below provides a quick overview of the differences and similarities between the OmniSwitch 9000 and OmniSwitch 9000E software and hardware features. Additional details are provided in the new hardware and new software sections.

Feature / Hardware Differences	OS-9000	OS-9000E
Authenticated VLANs	Supported	Not supported in 6.4.1
IPX	Supported	Not supported in 6.4.1
Binding Rules	Supported	Not supported in 6.4.1
Chassis and Fans	9600/9700/9800	Same 9700/9800 chassis (no support for 9600)
Power Supplies	Same Power Supplies	Same Power Supplies
CMMs	9600/9700/9800	9700E/9800E
24-port Copper Module	OS9-GNI-C24	OS9-GNI-C24E
24-port Universal Module	OS9-GNI-U24	OS9-GNI-U24E
2-port Gigabit Module	OS9-XNI-U2	OS9-XNI-U2E
Power Over Ethernet	Supported	Not supported in 6.4.1

System Requirements

Memory Requirements

OmniSwitch 9000E Series Release 6.4.1.R01 requires 1 GB of SDRAM and 256MB of flash memoryfor 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 9000E

Release	**UBoot**	**FPGA**	**Miniboot.uboot**
6.4.1.394.S01	6.4.1.149.R01	Major 0x2 Minor 0x19	6.4.1.149.R01

Note: Release 6.4.1 is only supported on an OmniSwitch 9000E. This release is not supported on any OmniSwitch 9000 CMMs or NIs.

New Hardware Supported

Note: OmniSwitch 9000E CMMs and NIs are supported in existing OmniSwitch 9700 and OmniSwitch 9800 chassis. However, OmniSwitch 9000 modules and OmniSwitch 9000E modules cannot be mixed in the same chassis.

Note: OmniSwitch 9000E modules **are not** supported in a OmniSwitch 9600 chassis.

OS9000E

The OmniSwitch 9000E chassis supports a high-performance switch fabric and provides slots for Ethernet, Gigabit Ethernet, and/or 10 Gigabit Ethernet Network Interface (NI) modules. An additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9000E supports redundant power supplies, 10/100/1000 copper ethernet ports, 1000 Mbps fiber ports as well as 10-Gigabit ethernet ports.

New Chassis Management Module (CMM)

The following CMMs are available in this release:

OS9700E-CMM / OS9800E-CMM

The Chassis Management Module (CMM) is the management unit for OmniSwitch 9000E switches. In its role as the management unit, the CMM provides key system services, including:

- Console, USB, and Ethernet management port connections to the switch
- Software and configuration management, including the Command Line Interface (CLI)
- Web-based management (WebView)
- SNMP management
- Power distribution
- Switch diagnostics
- Important availability features, including redundancy (when used in conjunction with another CMM), software rollback, temperature management, and power management
- The CMM also contains the switch fabric unit for the OmniSwitch 9000E. Data passing from one NI module to another passes through the CMM fabric. When two CMMs are installed, both fabrics are normally active.

Note. The USB port on the front panel of the CMM is not supported in this release.

Network Interface (NI) Modules

The following NI modules are available in this release:

OS9-GNI-C24E

Provides 24 auto-sensing twisted-pair ports, individually configurable as 10BaseT, 100BaseTX, or 1000BaseT.

OS9-GNI-U24E

Provides 24 SFP connectors for various gigabit, dual-speed, and bi-directional SFP transceivers.

OS9-XNI-U2E

Provides two 10-Gigabit XFP connectors for various XFP transceivers.

Power Supplies

OS9-PS-0600A

The PS-0600A is a 600 Watt AC power supply for OmniSwitch 9000E switches.

OS9-PS-0600D

The PS-0600D is a 600 Watt DC power supply for OmniSwitch 9000E switches.

SFP Transceivers - Gigabit Ethernet

SFP-GIG-CWD60

A group of 8 CWDM Gigabit Ethernet optical transceivers (SFP MSA). Supports single-mode fiber from 1470nm to 1610nm wavelength (based on transceiver) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.

SFP-GIG-EXTND

1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 850nm wavelength with an LC connector. Typical reach up to 2 Km on 62.5/125 µm MMF and 50/125 µm MMF.

SFP-GIG-LH40

1000Base-LH Gigabit Ethernet optical transceiver (SFP MSA). Supports single-mode fiber over 1310nm wavelength with an LC connector. Typical reach up to 40 Km on 9/125 µm SMF.

SFP-GIG-LH70

1000Base-LH Gigabit Ethernet optical transceiver (SFP MSA). Supports single-mode fiber over 1310nm wavelength with an LC connector. Typical reach up to 70 Km on 9/125 µm SMF.

SFP-GIG-LX

1000Base-LX Gigabit Ethernet optical transceiver (SFP MSA). Supports single-mode fiber over 1310nm wavelength with an LC connector. Typical reach up to 10 Km on 9/125 µm SMF.

SFP-GIG-SX

1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 850nm wavelength with an LC connector. Typical reach up to 300 m on 62.5/125 µm or 550m 50/125 µm MMF.

SFP-GIG-T

10/100/1000Base-T Gigabit Ethernet transceiver (SFP MSA). Supports category 5, 5E, and 6 copper cabling up to 100m.
***NOTE**: Only SFP-GIG-T transceivers shipped with an OS-9000E order support triple speed capability on an OS9000E. Other SFP-GIG-T transceivers support gigabit speed only.*

SFP-GIG-BX-D

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

SFP-GIG-BX-U

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

SFP Transceivers – Dual Speed Ethernet

SFP-DUAL-MM

Dual Speed 100Base-FX or 1000Base-LX Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 1310nm wavelength with an LC connector. Typical reach of 550m at Gigabit speed and 2Km at 100Mbit speed.

SFP-DUAL-SM10

Dual Speed 100Base-FX or 1000Base-LX Ethernet optical transceiver (SFP MSA). Supports single-mode fiber over 1310nm wavelength with an LC connector. Typical reach of 10Km at Gigabit speed and 100Mbit speed.

SFP Transceivers – 100 FX Ethernet

SFP-100-BX20LT

100Base-BX SFP transceiver with an SC type interface. This bi-directional transceiver is designed for use over single-mode fiber optic on a single strand link up to 20KM point-to-point. Transmits at 1550nm and receives at 1310nm optical signal.

SFP-100-BX20NU

100Base-BX SFP transceiver with an SC type interface. This bi-directional transceiver is designed for use over single-mode fiber optic on a single strand link up to 20KM point-to-point. Transmits at 1310nm and receives at 1550nm optical signal.

SFP-100-LC-MM

100Base-FX SFP transceiver with an LC type interface. This transceiver is designed for use over multimode fiber optic cable.

SFP-100-LC-SM15

100Base-FX SFP transceiver with an LC type interface. This transceiver is designed for use over single-mode fiber optic cable up to 15KM.

SFP-100-LC-SM40

100Base-FX SFP transceiver with an LC type interface. This transceiver is designed for use over single-mode fiber optic cable up to 40KM.

XFP Transceivers

XFP-10G-ER40

10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single-mode fiber over 1550nm wavelength with an LC connector. Typical reach of 40Km on 9/125 µm SMF.

XFP-10G-LR

10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single-mode fiber over 1310nm wavelength with an LC connector. Typical reach of 10Km on 9/125 µm SMF.

XFP-10G-SR

10 Gigabit Ethernet optical transceiver (XFP MSA). Supports multimode fiber over 850nm wavelength with an LC connector. Typical reach of 300m on 50/125 µm MMF.

XFP-10G-ZR80

10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single-mode fiber over 1550nm wavelength with an LC connector. Typical reach of 80Km on 9/125 µm SMF.

XFP-10G-CX4

10 Gigabit Ethernet copper transceiver (XFP MSA). Supports links up to 15 meters on standard CX4 type cables compliant with IEEE 802.3ak. InfiniBand type cables not supported.

New Software Features

The following software features are supported with the 6.4.1.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

Feature Summary

Feature	Platform	Software Package
31-bit Network Mask Support	9000E	base
802.1ab	9000E	base
802.1Q	9000E	base
802.1Q 2005 (MSTP)	9000E	base
802.1x Multiple Client Support	9000E	base
802.1x Device Classification (Access Guardian)	9000E	base
Access Control Lists (ACLs)	9000E	base
Access Control Lists (ACLs) for IPv6	9000E	base
Account & Password Policies	9000E	base
ACL & Layer 3 Security	9000E	base
ACL Manager (ACLMAN)	9000E	base
ARP Defense Optimization	9000E	base
ARP Poisoning Detect	9000E	base
Authenticated Switch Access	9000E	base
Auto-Qos Prioritization of NMS Traffic	9000E	base
Auto-Qos Prioritization of IP Phone Traffic	9000E	base
Automatic VLAN Containment (AVC)	9000E	base
BGP4	9000E	base advanced routing
BGP Graceful Restart	9000E	base advanced routing
BPDU Shutdown Ports	9000E	base
Command Line Interface (CLI)	9000E	base
DHCP Relay	9000E	base
DHCP Option-82	9000E	base
DHCP Snooping	9000E	base
DHCP Snooping Option-82 Data Insertion Format	9000E	base
DNS Client	9000E	base
DSCP Range Condition	9000E	base
Dynamic VLAN Assignment (Mobility)	9000E	base
DVMRP	9000E	base advanced routing
ECMP RIP Support	9000E	base
End User Partitioning	9000E	base
Ethernet Interfaces	9000E	base
Ethernet OAM	9000E	base
Flood/Storm Control	9000E	base

Generic Routing Encapsulation (GRE)	9000E	base
GVRP	9000E	base
Health Statistics	9000E	base
HTTP/HTTPS Port Configuration	9000E	base
Interswitch Protocols (AMAP)	9000E	base
IP-IP Tunneling	9000E	base
IPv4 Routing	9000E	base
IPv6 Routing	9000E	base
IP DoS Filtering	9000E	base
IPv4 Multicast Switching (IPMS)	9000E	base
IPv6 Multicast Switching (MLD)	9000E	base
IPv4 Multicast Switching (Proxying)	9000E	base
IP MC VLAN – Support for multiple sender ports	9000E	base
IPv6 Multicast Switching (Proxying)	9000E	base
IPv6 Client and/or Server Support	9000E	base
IPv6 Multicast Routing	9000E	advanced routing
IP Multinetting	9000E	base
IP Route Map Redistribution	9000E	base
IS-IS	9000E	advanced routing
ISSU	9000E	base
L2 DHCP Snooping	9000E	base
L2 Static Multicast Address	9000E	base
L4 ACLs over IPv6	9000E	base
Learned MAC Address Notificaton	9000E	base
Learned Port Security (LPS)	9000E	base
Link Aggregation (static & 802.3ad) – 128 Max Groups with maximum 2 ports per group	9000E	base
MAC Address Mode	9000E	base
Mac Authentication for Supplicant/non-Supplicant	9000E	base
NTP Client	9000E	base
OSPFv2	9000E	base advanced routing
OSPFv3	9000E	base advanced routing
Partitioned Switch Management	9000E	base
Per-VLAN DHCP Relay	9000E	base
PIM PIM-SSM (Source-Specific Multicast)	9000E	base advanced routing
Policy Server Management	9000E	base
Port-based Ingress Limiting	9000E	base
Policy Based Mirroring	9000E	base
Policy Based Routing (Permanent Mode)	9000E	base
Port Mapping	9000E	base

Port Mirroring (1:24)	9000E	base
Port Mirroring (1:128)	9000E	base
Port Monitoring	9000E	base
Power over Ethernet (PoE)	9000E	base
PVST+	9000E	base
Quality of Service (QoS)	9000E	base
Quarantine Manager and Remediation	9000E	base
Redirection Policies (Port and Link Aggregate)	9000E	base
Remote Port Mirroring	9000E	base
RIPv1/RIPv2	9000E	base
RIPng	9000E	base
RMON	9000E	base
Router Discovery Protocol (RDP)	9000E	base
Routing Protocol Preference	9000E	base
RRSTP	9000E	base
Secure Copy (SCP)	9000E	base
Secure Shell (SSH)	9000E	base
Server Load Balancing	9000E	base
SSH Public Key Authentication	9000E	base
sFlow	9000E	base
Smart Continuous Switching Hot Swap Management Module Failover Power Monitoring Redundancy	9000E	base
SNMP	9000E	base
Source Learning	9000E	base
Software Rollback	9000E	base
Spanning Tree	9000E	base
Syslog to Multiple Hosts	9000E	base
Switch Logging	9000E	base
Traffic Anomaly Detection (Network Security)	9000E	base
Text File Configuration	9000E	base
UDLD	9000E	base
User Definable Loopback Interface	9000E	base
User Network Profiles	9000E	base
VLANs	9000E	base
VLAN Stacking and Translation	9000E	base
VLAN Stacking Eservices	9000E	base
Multiple Virtual Routing & Forwarding (Multiple-VRF)	9000E	base
VRRPv2	9000E	base
VRRPv3	9000E	base
VRRP Global Commands	9000E	base
Web-Based Management (WebView)	9000E	base
Webview/SNMP support for BGP IPv6	9000E	advanced routing

Extensions		
Windows Vista for WebView	**9000E**	**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- supplicant 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-supplicant 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-supplicant 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.

The number of possible 802.1X users is 2K per system, not to exceed 1K per module. This number is a total number of users that applies to all authenticated clients, such as 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.

This release also supports 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.
- *Multicast ACLs*—for filtering IGMP traffic.

Access Control Lists (ACLs) for IPv6

Support for IPv6 ACLs on the OmniSwitch 9000E 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 **icmpcode**.
- **TCP connection rules**—Allows the determination of an *established* TCP connection by examining TCP flags found in the TCP header of the packet. Two condition parameters are available for defining a TCP connection ACL: **established** and **tcpflags**.
- **Early ARP discard**—ARP packets destined for other hosts are discarded to reduce processing overhead and exposure to ARP DoS attacks. No configuration is required to use this feature, it is always available and active on the switch. Note that ARPs intended for use by a local subnet and VRRP are *not* discarded.
- **UserPorts**—A port group that identifies its members as user ports to prevent spoofed IP traffic. When a port is configured as a member of this group, packets received on the port are dropped if they contain a source IP network address that does not match the IP subnet for the port.
- **UserPorts Profile**—In addition to spoofed traffic, it is also possible to configure a global UserPorts profile to specify additional types of traffic, such as BPDU, RIP, OSPF, DVMRP, PIM, 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.
- **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.

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 administator can configure how often a password must be changed, lockout settings for failed attempts, password complexity, history, and age as well as other account management settings.

ARP Defense Optimization

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

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.

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.

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, 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.

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 (GRE)

Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels. GRE tunnels on an OmniSwitch are used to create a wire-rate 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.

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 connctivity 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.

This release support 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
- OSPFv2
- BGP4
- Static Routes
- IP-MTU per VLAN – Allows the setting of the MTU for each IP interface associated with a VLAN.

  Note: The IP-MTU setting will affect both IPv4 and IPv6 interfaces.

  Note: Control plane traffic is not affected by the IP-MTU setting.

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 9000E 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 9000E. 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
- IP-MTU per VLAN – Allows the setting of the MTU for each IPv6 interface associated with a VLAN.

OmniSwitch 9000E 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
- Invalid IP Attack
- Multicast IP and MAC Address Mismatch
- Ping Overload
- Packets with loopback source IP address

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 9000E 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 9000E Series.

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.

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.

In-Service Software Upgrade (ISSU)

The In-Service Software Upgrade (ISSU) feature is used to patch the CMM images running on an OmniSwitch 9000E with minimal disruption to data traffic. The CMM images can be patched on a fully synchronized, certified, and redundant system running an ISSU capable build without requiring a reboot of the switch. Only non-NI related issues are ISSU capable.

- Switches running an '**R**##' build, such as 6.4.1.123.R01 **do not** support ISSU upgrades. The switch must first be upgraded to an '**S##**' build such as 6.4.1.123.**S01**.

- Periodic ISSU capable patches will be available on the Service & Support website. These patches will contain all CMM-only related fixes and will support the ISSU capability.

- ISSU patches are only supported within the same 'S##' branch. For example, if a switch is running 6.4.1.123.S01 then only 6.4.1.###.S01 images can used to perform an ISSU patch. If a switch is running 6.4.1.234.S02 then only 6.4.1.###.S02 images can used to perform an ISSU patch.

- Approximately every six months a new ISSU capable branch will be available from Service & Support. (i.e. S01, S02, S03, etc.). Each new branch will include all NI related fixes that were not supported in the previous ISSU branch. Upgrading from one ISSU branch to another will require a reboot and should be scheduled during a maintenance window.

- If a critical NI related patch is required, it will be necessary to move to an "**R**##" related build. Since "**R**##" related builds do not support the ISSU feature, a reboot will be required and should be scheduled during a maintenance window.

- The images which are ISSU capable are **Jbase.img**, **Jsecu.img**, **Jadvrout.img** and **Jos.img.**

- A minimum of 25 MB flash space must be present in the switch to accommodate the image files that are used to patch existing image files. This feature is only supported on the OmniSwitch 9000E.

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 link aggregation groups that can consist of 2, 4, or 8 Ethernet-ports with a maximum of 256 link aggregation ports and 128 link aggregation groups per switch.
  - 2 ports per group - maximum 128 link aggregate groups
  - 4 ports per group – maximum 64 link aggregate groups
  - 8 ports per group – maximum 32 link aggregate groups

- **Reliability**. If one of the physical links in a link aggregate group goes down, the link aggregate group can still operate.

Ease of Migration. Link aggregation can ease the transition from a Gigabit Ethernet backbone to a 10 Gigabit Ethernet backbone.

- **Interoperability with Legacy Switches**. Static link aggregation can interoperate with OmniChannel on legacy switches.

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).

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 9000E 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.

PVST+ Interoperability

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

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

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 9000E 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
- **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: BPDU Shutdown Ports group is *not* supported on the OmniSwitch 9000E 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-base mirroring are both allowed on the same port at the same time.
- One remote port-based mirroring session supported per switch.
- One port-monitoring session supported per switch.

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 9000E switches, configuring minimum and maximum egress bandwidth is supported on a per COS queue basis for each port

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 features makes use of an 802.1q tag to send the mirrored traffic over the network using tagged VLANs.

- There must not be any physical loop present in the remote port mirroring VLAN.
- Spanning Tree must be disabled for the remote port mirroring VLAN.
- BPDU mirroring will be enabled by default on all OS9000Es 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.
- 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.

OmniSwitch 9000E 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. Additionally, ECMP capability for up to 4 paths is also supported.

RIPng

The OmniSwitch 9000E 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).

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 9000E 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 9000E

Each OS9000E CMM module contains hardware and software elements to provide management functions for the OS9000E system. The OS9000E CMM module also contains the switch fabric for the OS9000E system. User data flowing from one NI module to another passes through the switch fabric.

The OS9000E will operate with one or two CMM modules installed.

If there are two CMM modules in an OS9000E, 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 OS9000E, both fabric modules are normally active. Two CMM modules must be installed in the switch 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 OS9000E, 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 9000E 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 9000E Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 32K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K per OmniSwitch 9000E chassis.

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. The switch supports up to 128 STP instances when running in 1X1 mode. 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. This 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.
- 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 Eservice Mode - The VLAN Stacking application operates in the Eservice mode only, legacy mode is not supported. Eservice mode offers the following enhancements:

- 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 Eservices mode consists of using an approach based on defining an Ethernet service to tunnel customer traffic.

Multiple Virtual Routing and Forwarding (Multiple-VRF)

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

Some of the benefits of using the Multiple VRF feature include the following:

- Multiple routing instances within the same physical switch. Each VRF instance is associated with a set of IP interfaces and creates and maintains independent routing tables. Traffic between IP interfaces is only routed and forwarded within those interfaces/routes that belong to the same VRF instance.
- Multiple instances of IP routing protocols, such as static, RIP, IPv4, BGPv4, and OSPFv2 on the same physical switch. An instance of each type of protocol operates within its own VRF instance.
- The ability to use duplicate IP addresses across VRF instances. Each VRF instance maintains its own IP address space to avoid any conflict with the service provider network or other customer networks.

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

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

Refer to the following table to determine the VRF association for a specific switch application. Applications that do not appear in this table are non-VRF aware.

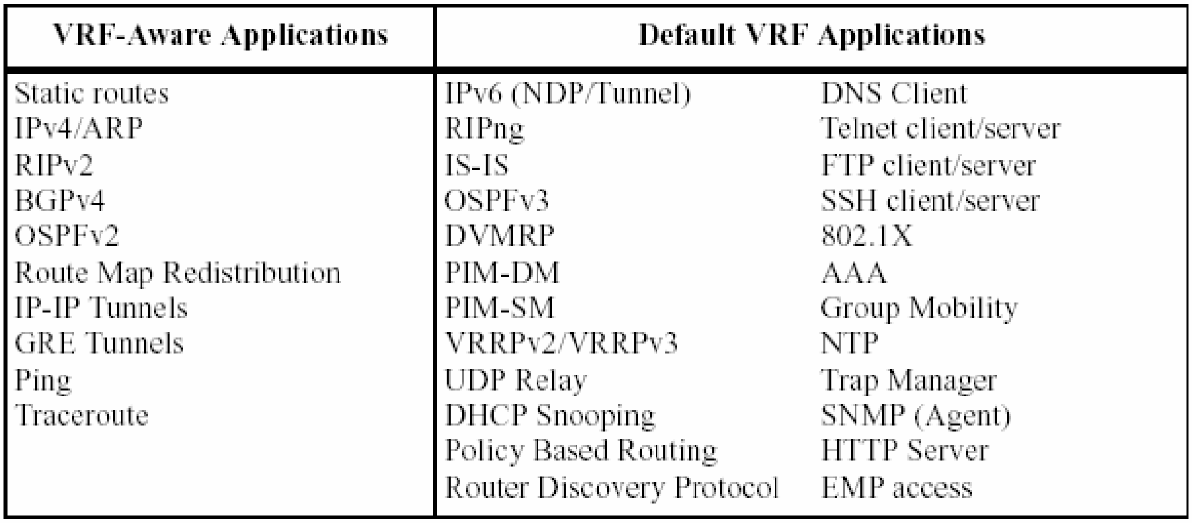

VRF-Aware Applications	Default VRF Applications	
Static routes	IPv6 (NDP/Tunnel)	DNS Client
IPv4/ARP	RIPng	Telnet client/server
RIPv2	IS-IS	FTP client/server
BGPv4	OSPFv3	SSH client/server
OSPFv2	DVMRP	802.1X
Route Map Redistribution	PIM-DM	AAA
IP-IP Tunnels	PIM-SM	Group Mobility
GRE Tunnels	VRRPv2/VRRPv3	NTP
Ping	UDP Relay	Trap Manager
Traceroute	DHCP Snooping	SNMP (Agent)
	Policy Based Routing	HTTP Server
	Router Discovery Protocol	EMP access

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

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.

Global VRRP Configuration

Includes 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 or 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 and 3.0 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 6.4.1.R01:

No.	Trap Name	Platforms	Description
0	coldStart	all	The SNMP agent in the switch is reinitiating and itsk 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 pesudo-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	n/a	Two or more slots claim to have the same slot number.
38	alaStackMgrNeighborChangeTrap	n/a	Indicates whether or not the stack is in loop.
39	alaStackMgrRoleChangeTrap	n/a	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	n/a	Occurs whenever a binding rule which has been configured gets violated.
43	unused	—	
44	unused	—	
45	unused	—	
46	unused	—	
47	pethPsePortOnOff	n/a	Indicates if power inline port is or is not delivering power to the a power

			inline device.
48	pethPsePortPowerMaintenanceStatus	n/a	Indicates the status of the power maintenance signature for inline power.
49	pethMainPowerUsageOn	n/a	Indicates that the power inline usage is above the threshold.
50	pethMainPowerUsageOff	n/a	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	n/a	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	n/a	The element identified by alaStackMgrSlotNINumber will enter the pass through mode because its operational slot was cleared with immediate effect.
56	alaStackMgrOutOfSlotsTrap	n/a	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	n/a	The element identified by alaStack MgrSlotNINumber will enter the pass through mode because there are no tokens available to be assigned to this element.
58	alaStackMgrOutOfPassThruSlotsTrap	n/a	There are no pass through slots avail able 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 configura tion.
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 monitor ing 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	n/a	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	n/a	The old PRIMARY element cannot be detected in the stack. There is a possiblity of a duplicate MAC address in the network
73	vRtrIsisDatabaseOverload	all	This notification is generated when the system enters or leaves the Overload state.
74	vRtrIsisManualAddressDrops	all	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
75	vRtrIsisCorruptedLSPDetected	all	This notification is generated when an LSP that was stored in memory has become corrupted.
76	vRtrIsisMaxSeqExceedAttempt	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 recieve 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 incorrent authentication information field.
83	vRtrIsisVersionSkew	all	A notification sent when a 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 Level2 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 adjancency'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
Authenticated VLANs	**all**	**base**
IPX	**all**	**base**
Binding Rules	**all**	**base**
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-eeprom** **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** **ip ospf redist-filter effect** **ip ospf redist-filter metric** **ip ospf redist-filter route-tag** **ip ospf redist-filter redist-control**
PIM	**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**
QoS	**qos classify fragments** **qos flow timeout** **show policy classify destination interface type** **show policy classify source interface type**
RIP	**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**
SYSTEM	**install**
VLANs	**vlan router mac multiple enable|disable** **vlan binding mac-port-protocol** **vlan binding mac-ip** **vlan binding ip-port**

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	

AlcatelIND1Port	esmPortCfgLongEnable esmPortCfgRuntEnable **esmPortCfgRuntSize** **esmPortPauseSlotTime** **esmPortCfgFLow**	**alcether10GigTable**
AlcatelIND1QoS	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 alaQoSConfigAppliedReflfexiveTimeout	

	alaQoSConfigFragmentTimeout alaQoSConfigAppliedFragmentTimeout alaQoSConfigClassifyFragments alaQoSConfigAppliedClassifyFragments	
AlcatelIND1Slb	slbFeature slbClusterTable **slbServerTableg**	
AlcatelIND1StackManager	**alaStackMgrStatsTable**	
AlcatelIND1SystemService	**systemUpdateStatusTable**	
AlcatelIND1VlanManager	vlanIpxNet vlanIpxEncap vlanIpxRipSapMode **vlanIpxDelayTicks** **vlanSetMultiRtrMacStatus**	vlanIpxStatus vlanSetIpxRouterCount
AlcatelIND1WebMgt	alaIND1WebMgtRFSConfigTable alaIND1WebMgtHttpPort **alaIND1WebMgtHttpsPort**	
IEEE_802_1X	dot1xAuthDiagTable dot1xAuthSessionStatsTable dot1xSuppConfigTable dot1xSuppStatsTable	
IETF_BGP4	**bgpRcvdPathAttrTable** **bgp** **bgpPeerTable** **bgp4PathAttrTabl**	
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 dot1dTrafficClassTable 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 sdlc 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

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.

113285	In WebView, if two or more Internet Explorer browser windows are opened for different routers and both sessions have any Add, Modify, or Help windows open simultaneously, the newly opened session may logout the previous session.	Two workarounds: 1) Close any Add, Modify, or Help windows in one session before opening any Add, Modify, or Help windows in another session. 2) Use only one browser window to configure a router at a time.
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 the Firefox browser with Windows Vista.
118563	Webview networking-IP interface page is unable to add an address to Loopback0.	Use the "ip interface" CLI command.
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.
119679	WV System->System Mgmt.-> File->Local backup creates mutiple files.	There is no known workaround at this time.
121160	WebView help is missing for the ISIS->Circuit Level->Configuration>Modify page.	There is no known workaround at this time.

124236	WebView fails to modify LLDP slot configuration.	Use the CLI to modify LLDP slot configuration.
125891	WebView DoS Attacks screen does not accurately reflect statistics.	Use the CLI to view Dos statistics.

SNMP

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 known workaround at this time.
119593	The PR refers to the READ Error being displayed, when any of the MIB objects - UdldGlobalConfigUdldProbeIntervalTimer, UdldGlobalConfigUdldDetectionPeriodTimer, UdldGlobalClearStats, UdldGlobalConfigUdldMode are read. The problem does not affect the UDLD functionality in any case. These MIB objects are basically significant for write purpose only and Get-Output of these objects do not stand valid.	There is no work around at this time. This has no affect on functionality.
123640	The SNMP authkey parameter returns an error and cannot be configured.	Authkey is a read-only parameter and cannot be configured.

CLI

117588	CLI will provide help for unsupported parameters when using '?'.	Refer to the CLI Reference guide for supported parameters.

LAYER 2

Ethernet

122798	Jabber frame counters do not get updated in the "show interfaces <s/p> accounting" command.	There is no known workaround at this time.
126116	On an OS9000E with an SFP-GIG-T transceiver, the duplex setting gets reset to full duplex after a reboot.	Re-configure the duplex setting for half duplex after a reboot.
127924	When changing the MTU setting for a VLAN with an IPv6 interface and then disabling and re-enabling the VLAN, the IPv6 interface will sometimes remain administratively disabled.	There is no known workaround at this time.

IPMV

126749	After extracting an NI, "write memory" does not clean up ipmvlan configuration on the extracted NI.	There is no known workaround at this time.

Spanning Tree

95308	Temporary traffic loops could happen under the following scenarios: 1. Reloading of a non root bridge. This happens when the bridge is going down and is due to the sequential bringing down of NIs during a reload process .It is purely temporary in nature and stops when all the NIs eventually get powered off. 2. NI power down When an NI power down command is executed for an NI and if that NI has the Root port port and other NIs have Alternate ports, it is possible to see some traffic looping back from the newly elected Root port. The traffic loop back is temporary and will stop once the NI gets powered off. 3. New Root bridge selection Temporary loops could occur during the process of electing a new Root bridge, if this election process is triggered by the assignment a worse priority for the existing root bridge or a root bridge failure. This happens due to the inconsistent spanning tree topology during the convergence and stops entirely once the network converges	For items 1 and 2 above there is no work around presently. For item 3 the following work around could be applied: 1. Tune the max age (and or max hops in the case of MSTP) parameter to a lower value that is optimal for the network. This will reduce the convergence time and thereby the duration of temporary loops. 2. To select a new root bridge, consider assigning better priority for that bridge instead of assigning worse priority for the existing root bridge.
105493	Enabling Spanning Tree in flat mode after disabling does not work when VLAN 1 is disabled.	Perform the following steps in the order of: 1. Enable VLAN 1 2. Enable STP 3. Disable VLAN 1
127343	RRSTP convergence time is sometimes greater than 50ms. This is only when the link is administratively brought down, not bridge down (power off).	Set vlan-tag for rrstp ring. ex) bridge rrstp ring 1 vlan-tag 20 This will raise the priority of RRSTP BPDU to provide moderate improvement.

VLAN Stacking

118736	Giving a CVLAN range bigger than 128, the sap sometimes gives error and inconsistent state.	There is no known workaround at this time.
125944	Can not delete cvlan from cli but uni can be deleted	Delete the sap profile and recreate it with new cvlan.

LAYER 3

General

109841	If filtering is used in command "show ip route", only one gateway will be displayed if there are multiple lines to display for ecmp routes.	There is no known workaround at this time..
120445	IP Multicast TTL Threshold capability is not currently supported.	There is no known workaround at this time.
127253	An OS9000E with MTU-IP at 9198 does not form OSPF adjacency with IPD switch which has max MTU set to 9194	There is no known workaround at this time.
127839	Graceful restart might not be successful in some scenarios of redistribution and multiple neighbors	There is no known workaround at this time.

IPv6

119506	IPv6 OSPFv3 cannot be configured by LDAP user with full read/write access.	Configure OSPFv3 with SNMP, CLI or Webview.
127921	With large MTU such as 9198, ospf3 can not form a FULL adjacency in a two switch setup.	There is no known workaround at this time.
128233	ping6 does not work when pinging the config-tunnel interfaces or 6to4 tunnel interfaces..	There is no known workaround at this time.

PIM-SM

119471	When configuring PIM, if the SSM range overlaps with the ASM range, some packets may be dropped.	Avoid overlapping SSM and ASM ranges.

VRF

123955	A Policy Based Routing rule configured in the default VRF instance may affect a flow in non-default VRF.	Policy Based Routing is part of QoS and QoS is not VRF aware. There is no known workaround at this time.
124169	IPRM allows the same interface name to be used in different route-maps in different VRFs.	The same interface name can be used in different route-maps in different VRFs. Interface names are associated with a specific VRF. Therefore, in the VRF that contains the interface. the route-map processing will match the interface. In the VRF that doesn't contain the interface, the route-map processing will fail to match a non-

		existent interface.
125096	Incorrectly typing a VRF instance name will create a new VRF instance with that name.	Since the system does not prompt before creating a new VRF instance, ensure the name is entered correctly.
126374	"show ip interface emp" fails to return on non-default VRF	There is no known workaround at this time.
127881	IP interfaces respond with incorrect default ip-ttl when in VRF mode.	There is no known workaround at this time.
128092	Deleting a VRF does not reset the BGP neighbor counter for that VRF.	There is no known workaround at this time.

Quality of Service

General

120592	If the SLB probe period is set to a value higher than the SLB timeout period, the probe value will get reset after a switch reboot.	Do not set the probe period to a value higher than the timeout period.
125104	If the ip address is in the non-default VRF then you cannot use the built-in policy network group Switch.	There is no known workaround at this time.
125157	IPV6 policies specifying ip protocol and/or other L4 attributes such as TCP/UDP src/dst port can match/effect only on IPV6 traffic having upto 2 next header specifications. IPV6 packet having more than two NH specifications, will not be matched or will be matched incorrectly.	There is no known workaround at this time.
122390	When adding a 5th MAC address using the 'policy mac group alaPhones' command, a warning message is being displayed instead of an error message.	There is no known workaround at this time.

Security

Traffic Anomaly Detection (Network Security)

91228	System does not detect IPv6 port scanning or other IPv6 denial of service attacks.	There is no known workaround at this time.

System

General

106811	When entering the 'show interface slot/port' command the "SFP/XFP" field output for a port having an SFP plugged in cannot differentiate between 100Fx and Bidirectional SFP & between Gigabit and CWDM SFP.	There is no known workaround at this time.
113928	A race condition sometimes occurs where a delete and add MAC message get to the CMM in the wrong order. The entry still exists correctly on the NI, and the Keep Alive eventually populates the entry on the NI with the protocol field of 0.	There is no known workaround at this time. This is a display issue only.

Chassis Supervision

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.
124938	The system will reboot if the secondary CMM is removed during the ISSU upgrade procedure.	Allow the ISSU upgrade procedure to complete before removing the secondary CMM.

NI System

127687	The port frame size gets reset to the default size (1553) after "no power ni/power ni" sequence.	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.