

Release Notes

OmniSwitch 6450

Release 6.6.2.R02

These release notes accompany release 6.6.2.R02 software for the OmniSwitch 6450. They provide important information on individual software and hardware features. 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.

General Release Information:

- This release is only supported on the OmniSwitch 6450-10.
- This release does not support stacking.
- This release does not support Metro features.

Contents

- **Related Documentation**, see page 3
- **System Requirements**, see page 4
- **Hardware Supported**, see page 5
- **Supported Software Features**, see page 6
- **SNMP Traps**, see page 39
- **Unsupported Software Features**, see page 46
- **Unsupported CLI Commands**, see page 46
- **Unsupported MIBs**, see page 47
- **Open Problem Reports, and Feature Exceptions**, see page 55
- **Technical Support**, see page 60

Related Documentation

These release notes should be used in conjunction with the OmniSwitch 6450 along with the associated manuals as listed below.

User manuals can be downloaded at:

<http://enterprise.alcatel-lucent.com/?dept=UserGuides&page=Portal>

- ***OmniSwitch 6450 Series Getting Started guide***
Describes the hardware and software procedures for getting an OmniSwitch 6450 Series switch up and running.
- ***OmniSwitch 6450 Series Hardware User Guide***
Complete technical specifications and procedures for all OmniSwitch 6450 Series chassis, power supplies, and fans.
- ***OmniSwitch 6450 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 6450 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 6450 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 6450 Transceivers Guide***
Includes SFP and XFP transceiver specifications and product compatibility information.
- ***Technical Tips, Field Notices, Upgrade Instructions***
Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

- OmniSwitch 6450 Series Release 6.6.2.R02 requires 256 MB of SDRAM and 128MB of flash memory. This is the standard configuration shipped.
- Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory. Use the **show hardware info** command to determine your SDRAM and flash memory.

OmniSwitch 6450 - Available Image Files

Image File	Base or Optional Software	Description
KFbase.img	Base	CMM Base
KFeni.img	Base	NI image for all Ethernet-type NIs
KFos.img	Base	CMM Operating System
KFsecu.img	Base	Security

New Hardware Supported

OmniSwitch 6450-10

Provides eight (8) 10/100/1000BaseT Ethernet ports, two (2) combo ports, two (2) non-combo SFP ports, and an internal AC power supply.

Note: Stacking is not supported in this release.

Transceivers - Supported

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 μ m or 550m 50/125 μ m MMF.

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 μ m SMF.

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 μ m 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 μ m SMF.

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-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-BXLC-D

100Base-BX SFP transceiver with an LC 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. This transceiver is normally used in the central office (OLT) transmits 1550nm and receives 1310nm optical signal. Designed for use with SFP-100-BXLC-U.

SFP-100-BXLC-U

100Base-BX SFP transceiver with an LC 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. This transceiver is normally used in the client (ONU) transmits 1310nm and receives 1550nm optical signal. Designed for use with SFP-100-BXLC-D.

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.

Transceivers - Not Supported**SFP-GIG-T**

10/100/1000Base-T Gigabit Ethernet transceiver (SFP MSA). Supports category 5, 5E, and 6 copper cabling up to 100m.

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 μ m SMF.

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.

Supported Hardware/Software Combinations

The following table shows the 6.6.2.R02 software releases that support each of the listed models:

Model	Part Number	6.6.2.R02	All Others Releases
6450-10	903005-90	Supported	Not Supported

Supported Feature Descriptions

The following software features were introduced with the 6.6.1.R01 and are supported in 6.6.2.R02 release:

Feature Summary

Feature	Platform	Software Package
802.1ab	OS6450	base
802.1Q	OS6450	base
802.1x Multiple Client Support	OS6450	base
802.1x Device Classification (Access Guardian)	OS6450	base
Mac Authentication for 802.1x Supplicants	OS6450	base
Access Guardian	OS6450	base
Captive Portal	OS6450	base
Captive Portal Web Pages	OS6450	base
Access Control Lists (ACLs)	OS6450	base
Access Control Lists (ACLs) for IPv6	OS6450	base
L4 ACLs over IPv6	OS6450	base
ACL & Layer 3 Security	OS6450	base
ARP Defense Optimization	OS6450	base
ARP Poisoning Detection	OS6450	base
Authenticated Switch Access	OS6450	base
Partitioned Switch Management	OS6450	base
Account & Password Policies	OS6450	base
Command Line Interface (CLI)	OS6450	base
DHCP Relay	OS6450	base
Per-VLAN DHCP Relay		
DHCP Option-82	OS6450	base
DHCP Snooping	OS6450	base
L2 DHCP Snooping	OS6450	base
Option-82 Data Insertion Format	OS6450	base
DNS Client	OS6450	base
Dynamic VLAN Assignment (Mobility)	OS6450	base
End User Partitioning	OS6450	base
Ethernet Interfaces	OS6450	base
Flood/Storm Control	OS6450	base
Flow Control (802.3x)	OS6450	base
GVRP	OS6450	base
Health Statistics	OS6450	base
HTTP/HTTPS Port Configuration	OS6450	base
Interswitch Protocols (AMAP)	OS6450	base
IPv4 Routing	OS6450	base
31-bit Network Mask Support	OS6450	base
IPv6 Routing	OS6450	base
IPv6 Client and/or Server Support	OS6450	base
IP DoS Filtering	OS6450	base
IPv4 Multicast Switching (IPMS)	OS6450	base

Feature	Platform	Software Package
IPv6 Multicast Switching (MLD)	OS6450	base
IPv4 Multicast Switching (Proxying)	OS6450	base
IPv6 Multicast Switching (Proxying)	OS6450	base
IP MC VLAN (Multiple Sender Ports)	OS6450	base
IP Multinetting	OS6450	base
IP Route Map Redistribution	OS6450	base
Learned Port Security (LPS)	OS6450	base
Learned MAC Address Notificaton	OS6450	base
Link Aggregation (static & 802.3ad)	OS6450	base
Mac Retention	OS6450	base
NTP Client	OS6450	base
Policy Server Management	OS6450	base
Policy Based Routing (Permanent Mode)	OS6450	base
Port Mapping	OS6450	base
Port Mirroring (24:1)	OS6450	base
Port Monitoring	OS6450	base
Quality of Service (QoS)	OS6450	base
Auto-Qos Prioritization of IP Phone Traffic	OS6450	base
Auto-Qos Prioritization of NMS Traffic	OS6450	base
DSCP Range Condition	OS6450	base
Policy Based Mirroring	OS6450	base
Port-based Ingress Limiting	OS6450	base
Redirection Policies (Port and Link Agg)	OS6450	base
Tri-Color Marking	OS6450	base
Remote Port Mirroring	OS6450	base
RIPv1/RIPv2	OS6450	base
ECMP RIP Support	OS6450	base
RIPng	OS6450	base
RMON	OS6450	base
Router Discovery Protocol (RDP)	OS6450	base
Routing Protocol Preference	OS6450	base
Secure Copy (SCP)	OS6450	base
Secure Shell (SSH)	OS6450	base
SSH Public Key Authentication	OS6450	base
sFlow	OS6450	base
SNMP	OS6450	base
Source Learning	OS6450	base
- L2 Static Multicast Address	OS6450	base
Spanning Tree	OS6450	base
802.1Q 2005 (MSTP)	OS6450	base
Automatic VLAN Containment (AVC)	OS6450	base
PVST+	OS6450	base
RRSTP	OS6450	base
Switch Logging	OS6450	base
Syslog to Multiple Hosts	OS6450	base
Trivial File Transfer Protocol (TFTP) Client	OS6450	base

Feature	Platform	Software Package
Text File Configuration	OS6450	base
User Definable Loopback Interface	OS6450	base
User Network Profiles	OS6450	base
VLANs	OS6450	base
Web-Based Management (WebView)	OS6450	base

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, such as OmniVista, that deals with topology discovery. Switches that are compliant with 802.1AB use TLV (Time, Length, Value) frames to exchange information with neighboring devices and maintain a database of the information exchanged. The Link Layer Discovery Protocol-Media Endpoint Discover (LLDP-MED) is designed to extend IEEE 802.1AB functionality to exchange information such as VLANs and power capabilities.

802.1Q

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.

Access Guardian

802.1X Device Classification Policies

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 256 users per NI. 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.

Classification of both supplicant and non-supplicant devices using non-supplicant device classification policies is supported. As a result, MAC authentication is now applicable to both supplicant and non-supplicant devices.

Captive Portal

Captive Portal authentication is a configurable option within Access Guardian that allows Web browser clients to authenticate through the switch using 802.1x or MAC authentication via a RADIUS server. When the Captive Portal option is invoked, a Web page is presented to the user device to prompt the user to enter login credentials. If authentication returns a VLAN ID, the device is assigned to that VLAN. If a VLAN ID

is not returned or authentication fails, a separate Captive Portal policy then determines the network access control for the supplicant or non-supplicant.

Captive Portal Web Pages

Customizing the following Captive Portal Web page components is allowed. These components are incorporated and displayed when the Web-based login page is presented to the user.

- Logo
- Welcome text
- Background image
- User Acceptable Policy text
- Login help page

Captive Portal checks the local switch for any customized files before presenting the login Web page to the user. If any such files exist, they are incorporated into the Web page display. If no such files exist, the default Web page components are used.

Captive Portal Browser Support

The Captive Portal authentication feature presents the user with a Web page for entering login credentials. The following table provides the platforms and browser support information for Captive Portal users:

Platforms Supported	Web Browser Supported
Windows XP	IE6, IE7, IE8, FireFox2 and FireFox3
Windows Vista	IE7, Firefox2 and Firefox3
Linux	Firefox2 and Firefox3

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. Typically 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

The following QoS policy conditions are available for configuring ACLs to filter IPv6 traffic:

```

source ipv6
destination ipv6
ipv6
  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.

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 or VRRP are *not* discarded.
- **UserPorts**—A port group that identifies its members as user ports to prevent spoofed IP traffic. When a port is configured as a member of this group, packets received on the port are dropped if they contain a source IP network address that does not match the IP subnet for the port.
- **UserPorts Profile**—In addition to spoofed traffic, it is also possible to configure a global UserPorts profile to specify additional types of traffic, such as BPDU, RIP, OSPF, DVMRP, PIM, DHCP server response packets and DNS, 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.

ARP Defense Optimization

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

ARP Poisoning Detection

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

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

Authenticated Switch Access

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

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

- Remote Authentication Dial-In User Service (RADIUS). Authentication using this type of server was certified with 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.

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*.

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

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.

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.

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:

- **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.

DHCP Snooping – Layer 2

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 Omnsnitch 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.

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.

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.

Ethernet Interfaces

Ethernet and Gigabit Ethernet port software is responsible for a variety of functions that support Ethernet and Gigabit Ethernet. 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).

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 specific VLAN on the switch.

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.

Using 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. A trap will be sent if the number of dynamic VLANs exceeds the maximum threshold configured for GVRP.

Health Statistics

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

Health Monitoring provides the following data to the NMS:

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

- For each monitored resource, the following variables are defined:
- Most recent utilization level (percentage)
- Average utilization level over the last minute (percentage)
- Average utilization level over the last hour (percentage)
- Maximum utilization level over the last hour (percentage)
- Threshold level

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

HTTP/HTTPS Port Configuration

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

IP Multicast VLAN

The IP Multicast VLAN feature provides the ability to configure specific VLANs that are dedicated to distributing multicast traffic. These distribution VLANs connect to the nearest multicast router and support multicast traffic only.

IP Multicast VLANs are supported 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. Multiple sender ports are supported.

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
 - Ping
 - Traceroute
 - Address Resolution Protocol (ARP)
 - Internet Control Message Protocol (ICMP)
 - RIP I / RIP II
 - ECMP
 - Static routes

The base IP software allows you 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 you to trace an IP route, display Transmission Control Protocol (TCP) information, and display User Datagram Protocol (UDP) information.

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

31-Bit Network Mask Support – Configuring a 31-bit netmask is supported to allow for a point-to-point Ethernet network between two routers.

IPv6 Support

IPv6 (documented in RFC 2460) is designed as a successor to IPv4 and is supported on the OmniSwitch. 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
- RIPng
- Static Routes
- Ping6

- Traceroute6
- DNS client using Authority records
- Telnetv6 - Client and server
- File Transfer Protocol (FTPv6) – Client and server
- SSHv6 – Client and Server

The OmniSwitch supports hardware-based IPv6 routing.

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 an OmniSwitch 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.

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.

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.

Note that LPS is not configurable on link aggregate ports.

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

Link Aggregation (static & 802.3ad)

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

- **Scalability.** You can configure up to 32 link aggregation groups that can consist of 2, 4, or 8 Ethernet-ports.

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

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).

Policy Server Management

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

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

Port Mapping (Private VLANs)

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

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

Port Mirroring

When Port Mirroring is enabled, the active "mirrored" port transmits and receives network traffic normally, and the "mirroring" port receives a copy of all transmit and receive traffic to the active port. You can connect an RMON probe or network analysis device to the mirroring port to see an exact duplication of traffic on the mirrored port without disrupting network traffic to and from the mirrored port.

Up to two Port Mirroring sessions are supported per switch, one of which can be an RSPAN session. The session can be configured to a “N-to-1” session, where up to 24 source ports can be mirrored to a single destination port.

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.

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.

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 or denying traffic, or as complicated as remapping 802.1p bits from a Layer 2 network to ToS values in a Layer 3 network. OmniSwitch 6450 switches support 8 queues per port.

QoS is implemented on the switch through the use of policies, created on the switch or stored in Policy-View. 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.
- **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.

Note: If a policy rule only contains Layer 2 conditions (source MAC, destination MAC, VLAN, 802.1p, source port), then the rule is only applied to Layer 2 traffic. To apply a pure Layer 2 rule to Layer 3 traffic, add the “source ip any” keywords to a condition for that rule. To apply a pure Layer 2 rule to IPv6 traffic, add the “ipv6” keyword to a condition for that rule.

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.

DSCP Ranges – Configuring a range of DSCP values in a single QoS DSCP policy condition is now supported. This eliminates the need for multiple condition statements to configure multiple DSCP values for traffic classification. In addition, specifying a mask value is no longer required; QoS automatically calculates the appropriate mask value for each DSCP value specified.

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 from a source address.
- Traffic to a destination address.
- Traffic to and from an address.
- Traffic between 2 addresses.
- Traffic with a classification criterion based on packet contents other than addresses (for example , based on protocol, priority).
- VLAN-based mirroring - mirroring of packets entering a VLAN.

Policy-Based Mirroring limitations:

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

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.

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. On the OmniSwitch, configuring maximum egress bandwidth is supported on a per COS queue basis for each port

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

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

The TCM Meter operates in Color-Blind mode (the Color-Aware mode is not supported). In the Color-Blind mode, the Meter assumes that the incoming packet stream is uncolored.

There are two types of TCM marking supported:

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

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

Remote Port Mirroring (802.1Q Based)

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

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

The OmniSwitch supports RIP version 1 (RIPv1), RIP version 2 (RIPv2), and RIPv2 that is compatible with RIPv1. In addition, text key and MD5 authentication, on an interface basis, for RIPv2 is also supported. ECMP capability for up to 4 paths is also supported.

RIPng

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

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.

Note: The ingress and egress ports that participate in redirection policies must belong to the same VLAN.

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.

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.

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).

SNMP

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

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.

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.

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.

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.

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.

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.

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

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.

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 when a link failure occurs. 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).

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.

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

Trivial File Transfer Protocol (TFTP) Client

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

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.

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 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.

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, Vista
- Firefox 2.0 and later for Windows and Solaris SunOS 5.10

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

SNMP Traps

The following traps are supported in 6.6.2.R02 on the noted platforms:

No.	Trap Name	Platforms	Description
0	coldStart	all	The SNMP agent in the switch is reinitiating and its configuration may have been altered.
1	warmStart	all	The SNMP agent in the switch is reinitiating itself and its configuration is unaltered.
2	linkDown	all	The SNMP agent in the switch recognizes a failure in one of the communications links configured for the switch.
3	linkUp	all	The SNMP agent in the switch recognizes that one of the communications links configured for the switch has come up.
4	authenticationFailure	all	The SNMP agent in the switch has received a protocol message that is not properly authenticated.
5	entConfigChange	all	An entConfigChange notification is generated when a conceptual row is created, modified, or deleted in one of the entity tables.
6	aipAMAPStatusTrap	all	The status of the Alcatel-Lucent Mapping Adjacency Protocol (AMAP) port changed.
7	aipGMAPConflictTrap	—	This trap is not supported.
8	policyEventNotification	all	The switch notifies the NMS when a significant event happens that involves the policy manager.
9	chassisTrapsStr	all	A software trouble report (STR) was sent by an application encountering a problem during its execution.
10	chassisTrapsAlert	all	A notification that some change has occurred in the chassis.
11	chassisTrapsStateChange	all	An NI status change was detected.
12	chassisTrapsMacOverlap	all	A MAC range overlap was found in the backplane eeprom.
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.
20	esmDrvTrapDropsLink	all	This trap is sent when the Ethernet code drops the link because of excessive errors.

No.	Trap Name	Platforms	Description
21	pimNeighborLoss	all	This trap is not supported.
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	sIPCAMStatusTrap	all	The trap status of the Layer 2 pseudo-CAM for this NI.
31	unused	—	
32	unused	—	
34	ifMauJabberTrap	all	This trap is sent whenever a managed interface MAU enters the jabber state.
35	sessionAuthenticationTrap	all	An authentication failure trap is sent each time a user authentication is refused.
36	trapAbsorptionTrap	all	The absorption trap is sent when a trap has been absorbed at least once.
37	alaStackMgrDuplicateSlotTrap	-	Two or more slots claim to have the same slot number.
38	alaStackMgrNeighborChangeTrap	-	Indicates whether or not the stack is in loop.
39	alaStackMgrRoleChangeTrap	-	Indicates that a new primary or secondary stack is elected.
40	lpsViolationTrap	all	A Learned Port Security (LPS) viola-

No.	Trap Name	Platforms	Description
			tion has occurred.
41	alaDoSTrap	all	Indicates that the sending agent has received a Denial of Service (DoS) attack.
42	gmBindRuleViolation	all	Occurs whenever a binding rule which has been configured gets violated.
43	unused	—	
44	unused	—	
45	unused	—	
46	unused	—	
47	pethPsePortOnOff	-	Indicates if power inline port is or is not delivering power to the a power inline device.
48	pethPsePortPowerMaintenanceStatus	-	Indicates the status of the power maintenance signature for inline power.
49	pethMainPowerUsageOn	-	Indicates that the power inline usage is above the threshold.
50	pethMainPowerUsageOff	-	Indicates that the power inline usage is below the threshold.
53	httpServerDoSAttackTrap	all	This trap is sent to management station(s) when the HTTP server is under Denial of Service attack. The HTTP and HTTPS connections are sampled at a 15 second interval. This trap is sent every 1 minute while the HTTP server detects it is under attack.
54	alaStackMgrDuplicateRoleTrap	-	The element identified by alaStackMgrSlotNINumber detected the presence of two elements with the same primary or secondary role as specified by alaStackMgrChasRole on the stack.
55	alaStackMgrClearedSlotTrap	-	The element identified by alaStackMgrSlotNINumber will enter the pass through mode because its operational slot was cleared with immediate effect.
56	alaStackMgrOutOfSlotsTrap	-	One element of the stack will enter the pass through mode because there are no slot numbers available to be assigned to this element.
57	alaStackMgrOutOfTokensTrap	-	The element identified by alaStackMgrSlotNINumber will enter the pass through mode because there are no tokens available to be assigned to this element.
58	alaStackMgrOutOfPassThruSlotsTrap	-	There are no pass through slots available to be assigned to an element that is supposed to enter the pass through

No.	Trap Name	Platforms	Description
59	gmHwVlanRuleTableOverloadAlert	all	mode. An overload trap occurs whenever a new entry to the hardware VLAN rule table gets dropped due to the overload of the table.
60	lnkaggAggUp	all	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.
61	lnkaggAggDown	all	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
62	lnkaggPortJoin	all	This trap is sent when any given port of the link aggregate group goes to the attached state.
63	lnkaggPortLeave	all	This trap is sent when any given port detaches from the link aggregate group.
64	lnkaggPortRemove	all	This trap is sent when any given port of the link aggregate group is removed due to an invalid configuration.
65	pktDrop	all	The pktDrop trap indicates that the sending agent has dropped certain packets (to blocked IP ports, from spoofed addresses, etc.).
66	monitorFileWritten	all	A File Written Trap is sent when the amount of data requested by the user has been written by the port monitoring instance.
69	gmHwMixModeSubnetRuleTableOverloadAlert	all	A subnet overload trap occurs in mixed mode whenever a new entry to the HW subnet rule table gets dropped in OS6800 due to the overload of the table.
70	pethPwrSupplyConflict	-	Power supply type conflict trap.
71	pethPwrSupplyNotSupported	-	Power supply not supported trap.
72	lpsPortUpAfterLearningWindowExpiredTrap	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. This trap will also be generated at the time the Learning Window expires, with a slice and port value of 0.
92	dot1agCfmFaultAlarm	all	A MEP has lost contact with one or more MEPs. A notification (fault alarm) is sent to the management entity

No.	Trap Name	Platforms	Description
			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	chassisTrapsPossibleDuplicateMac	all	The old PRIMARY element cannot be detected in the stack. There is a possibility of a duplicate MAC address in the network.
101	lpsLearnMac	all	Generated when an LPS port learns a bridged MAC address.
102	gvrpVlanLimitReachedEvent	all	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
105	udldStateChange	-	Generated when the state of the UDLD protocol changes.
106	healthMonIpcTrap		IPC pools exceed usage/ causing trap."
107	Reserved	-	-
108	Reserved	-	-
109	arpMaxLimitReached	all	Generated when the hardware table has reached supported maximum entries.
110	ndpMaxLimitReached	all	Generated when the hardware table has reached supported maximum entries.
111	ripRouteMaxLimitReached	all	Generated when RIP database has reached supported maximum entries. RIP will discard any new updates.
112	ripngRouteMaxLimitReached	all	Generated when RIPng database has reached supported maximum entries. RIPng will discard any new updates.
113- 118	Reserved	-	
119	dot3OamThresholdEvent	-	This trap is sent when a local or remote threshold crossing event is detected. A local threshold crossing event is detected by the local entity, while a remote threshold crossing event is detected by the reception of an Ethernet OAM Event Notification OAMPDU that indicates a threshold event.
120	dot3OamNonThresholdEvent	-	This trap is sent when a local or remote non-threshold crossing event is detected. A local event is detected by

No.	Trap Name	Platforms	Description
			the local entity, while a remote event is detected by the reception of an Ethernet OAM Event Notification OAMPDU that indicates a non-threshold crossing event.
121	alaDot3OamThresholdEventClear	-	This trap is sent when is sent when a local or remote threshold crossing event is recovered.
122	alaDot3OamNonThresholdEventClear	-	This trap is sent is sent when a local or remote non-threshold crossing event is recovered.
123-	Reserved	-	
146			
147	halHashCollisionTrap	all	This trap is sent when an SFP/XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
148	alaLbdStateChangeToShutdown	all	This trap is sent when an SFP/XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
149	alaLbdStateChangeForClearViolationA	all	This trap is sent when an SFP/XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
150	alaLbdStateChangeForAutoRecovery	all	This trap is sent when an SFP/XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
151	Reserved	-	Reserved
152	Reserved	-	Reserved
153	alaErpRingStateChanged	-	This trap is sent when the ERP Ring State has changed.
154-	Reserved	-	Reserved
158			
159	alaDhcpClientAddressAddTrap	-	This trap is sent when a new IP address

No.	Trap Name	Platforms	Description
			is assigned to DHCP Client interface.
160	alaDhcpClientAddressExpiryTrap	-	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address
161	alaDhcpClientAddressModifyTrap	-	This trap is sent when the DHCP client is unable to obtain the existing IP address and a new IP address is assigned to the DHCP client.
162	alaDyingGaspTrap	-	This trap is sent when a switch has lost all power.
163	alaTestOamTxDoneTrap	-	After a configured time interval, this trap is sent to the NMS from Generator switch when the test duration expires.
164	alaTestOamRxReadyTrap	-	This trap is sent to the NMS once the switch with Analyzer or Loopback Role is ready to receive test traffic. Once this trap is received, the Generator is activated for generating test traffic.
165	alaTestOamTestAbortTrap	-	This trap is sent to the NMS from the switch, if the test is aborted during takeover.
166	Reserved	-	-
167	Reserved	-	-
168	alaSaaPIterationCompleteTrap	-	-
169	alaSaaEthIterationCompleteTrap	-	-
170	alaSaaMacIterationCompleteTrap	-	-

Unsupported Features

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

Feature	Platform	Software Package
CPE Testhead	OS6450	Base
Ethernet Ring Protection	OS6450	Base
Ethernet OAM	OS6450	Base
Ethernet Services	OS6450	Base
Service Assurance Agent	OS6450	Base
Mac Retention	OS6450	Base
Stacking	OS6450	Base

Unsupported CLI Commands

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

Software Feature	Unsupported CLI Commands
Chassis Mac Server	mac-range local mac-range duplicate-EEPROM mac-range allocate-local-only show mac-range status
DHCP Relay	ip helper traffic-suppression ip helper dhcp-snooping port traffic-suppression
Flow Control	flow
Hot Swap	reload ni [slot] # [no] power ni all
Interfaces	show interface slot/port hybrid copper counter errors show interface slot/port hybrid fiber counter errors
QoS	qos classify fragments qos flow timeout
System	install power ni [slot]

Unsupported MIBs

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

Feature	MIB
BGP	AlcatelIND1Bgp IETF_BGP4
DVMRP	AlcatelIND1Dvmrp IETF_DVMRP_STD_DRAFT
IPSec	AlcatelIND1IPsec.mib
IS-IS	AlcatelIND1Isis IETF_ISIS
Multicast Routing	AlcatelIND1Ipmm AlcatelIND1IpMcastDraft
OSPF/OSPFv3	AlcatelIND1DrcTm AlcatelIND1Ospf AlcatelIND1Ospf3 IETF_OSPF IETF_OSPFv3 IETF_OSPF_TRAP
PIM	AlcatelIND1Pim AlcatelIND1PimBsrDraft AlcatelIND1PimStdDraft IETF_PIM
Quality of Service (QoS)	IETF_P_BRIDGE
SLB	AlcatelIND1Slb.mib
Traffic Anomaly Detection	AlcatelIND1Ns
VRRP	AlcatelIND1VRRP.mib AlcatelIND1VRRP3.mib

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

WebView

PR	Description	Workaround
161256	WebView does not display success or failure status when performing a file transfer using TFTP.	There is no known workaround at this time.

Multicast

General

PR	Description	Workaround
160694	New multicast source does not get registered when the switch has already registered a different source for the same group on the same port.	There is no known workaround at this time.
161666	Multicast traffic is flooded across all the ports of a VLAN even when a static multicast member port is configured.	There is no known workaround at this time.
161813	For MLDv1 and MLDv2 a maximum of 1024 IPv6 multicast groups can be learned at the rate of 500pps.	There is no known workaround at this time.
162282	The number of groups forwarded to querier by IPMS enabled switch is less than what the switch has learned.	There is no known workaround at this time.

Quality of Service

General

PR	Description	Workaround
161371	If a QoS redirect policy is applied with a generic condition such as 'policy condition c1 source vlan 100' it not only redirects the	Create a QoS redirect policy with a more specific condition such as 'policy condition c1 source vlan 100

data traffic on VLAN 100 but also the L2 control plane packets. This could cause issues with protocols such as STP.

source ip any' to redirect IP packets but not control plane packets.

162340	Maximum QoS rules supported should display 1024 instead of 1280 on 'show qos statistics' command.	There is no known workaround at this time.
--------	---	--

Security

Captive Portal

PR	Description	Workaround
161944	The Captive Portal logout attempt fails if trying to logout after 3 to 5 minutes of login when using Netscape, Firefox or Chrome browsers.	Disconnect and reconnect the supplicant/non-supplicant user and flush out the entry from the switch.
162034	FTP and Telnet session gets disconnected when the supplicant is classified with the CP guest VLAN.	There is no known workaround at this time.

System

General

PR	Description	Workaround
160265	An error similar to "FfxVbf:CopyAndDiscardLguData() encountered a corrected data error in one or more of the pages" may be seen when issuing the 'write memory' command.	This is an informational file system error indicating a bit error was found and corrected. It has no affect on the system.
161369	When interface speed is changed to 100 mbps, the maximum frame size is set to 1553 even if a different frame size has been configured.	After changing the speed reconfigure the desired frame size.

Port Mapping

PR	Description	Workaround
159959	When performing many-to-one port mirroring some packets will not be mirrored if they are received by the mirroring port simultaneously or are being received at greater than wire speed.	There is no known workaround at this time.

Hot Swap / Redundancy

Feature Exceptions

CMM and Power Redundancy Feature Exceptions for OmniSwitch

- Stacking is NOT supported in 6.6.2.R02.
- The OmniSwitch 6450-10 does not support a backup power supply.

Technical Support

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

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

Email: esd.support@alcatel-lucent.com

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

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

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

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

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

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