Release Notes

OmniSwitch 6800/6850/9000

Release 6.1.5.R01

These release notes accompany release 6.1.5.R01 software for the OmniSwitch 6800, 6850, and 9000 hardware. They provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

Contents

Related Documentation

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

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

- *OmniSwitch 6800 Series Getting Started guide*

  Describes the hardware and software procedures for getting an OmniSwitch 6800 Series switch up and running.

- *OmniSwitch 6850 Series Getting Started Guide*

  Describes the hardware and software procedures for getting an OmniSwitch 6850 Series switch up and running.

- *OmniSwitch 9000 Series Getting Started Guide*

  Describes the hardware and software procedures for getting an OmniSwitch 9000 Series switch up and running.

- *OmniSwitch 6800 Series Hardware User Guide*

  Complete technical specifications and procedures for all OmniSwitch 6800 Series chassis, power supplies, and fans.

- *OmniSwitch 6850 Series Hardware User Guide*

  Complete technical specifications and procedures for all OmniSwitch 6850 Series chassis, power supplies, and fans.

- *OmniSwitch 9000 Series Hardware User Guide*

  Complete technical specifications and procedures for all OmniSwitch 9000 Series chassis, power supplies, and fans.

- *OmniSwitch CLI Reference Guide*

  Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.

- *OmniSwitch 6800/6850/9000 Network Configuration Guide*

  Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethernet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.

- *OmniSwitch 6800/6850/9000 Series Switch Management Guide*

  Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).

- *OmniSwitch 6800/6850/9000 Series Advanced Routing Configuration Guide*

  Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, and OSPFv3.

- *Upgrade Instructions for 6.1.5.R01*

  Provides instructions for upgrading the OmniSwitch 6800, 6850, 9000 to 6.1.5.R01.

- *OmniSwitch Transceivers Guide*

  Includes SFP and XFP transceiver specifications and product compatibility information.

- Technical Tips, Field Notices

Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

- OmniSwitch 6800 Series Release 6.1.5.R01 requires 256 MB of SDRAM and 64MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 6850 Series Release 6.1.5.R01 requires 256 MB of SDRAM and 64MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 9000 Series Release 6.1.5.R01 requires 256 MB of SDRAM and 128MB of flash memory for the Chassis Management Module (CMM). This is the standard configuration shipped.

Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory. Use the **show hardware info** command to determine your SDRAM and flash memory.

UBoot, FPGA, Miniboot, BootROM, and Upgrade Requirements

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

OmniSwitch 6800 Series

- Miniboot: 6.1.2.261.R03
- BootROM: 6.1.2.261.R03

OmniSwitch 6850 Series

- UBoot: 6.1.3.601.R01
- Miniboot.uboot: 6.1.3.601.R01

OmniSwitch 9000 Series

- UBoot NI: 6.1.1.167.R02; 6.1.5.354.R01 recommended.
- UBoot CMM: 6.1.1.167.R02; 6.1.5.354.R01 recommended.
- Miniboot.uboot CMM: 6.1.1.167.R02; 6.1.5.354.R01 recommended.
- FPGA CMM: Major Revision: 2 Minor Revision: 25 (displays as 0x19)

POE Firmware

- 5.01

Note. Refer to the *Upgrading OmniSwitch 6800, 6850, and 9000 Series Switches to 6.1.5.R01* instructions if a switch upgrade is necessary to meet the above requirements.

New Hardware Supported

The following new hardware is supported subject to the feature exceptions and problem reports described later in these release notes.

New Network Interface (NI) Modules

The following NI modules are available in this release:

OS9-GNI-C48T

Provides 48 auto-sensing ports using MRJ-21 connectors. Ports are auto-negotiating *and* individually configurable as 10BaseT, 100BaseTX, or 1000BaseT. To provide RJ45 connectivity, a special cable and patch panel is used to connect OS9-GNI-C48T ports to an RJ-45 port patch panel.

MRJ-21 to RJ45 Cable

The MRJ-21 cable is used to distribute a group of 6 RJ45 ports from the MRJ-21 connector on the OS9-GNI-C48T module to the RJ45 to RJ45 Patch Panel. The MRJ-21 cable is available in both 1 and 3 meter lengths.

RJ45 to RJ45 Patch Panel

The RJ45 to RJ45 Patch Panel contains 48 RJ45 connectors on both the front and back of the panel. It is used in conjunction with the MRJ-21 to RJ45 cable to allow for the distribution of 8 MRJ-21 connectors, each containing 6 Ethernet ports, from the OS9-GNI-C48T.

OS9-GNI-C20L

Provides 20 auto-sensing twisted-pair ports plus 2 SFP connectors. The 20 copper ports are auto-negotiating and individually configurable as 10BaseT or 100BaseTX, but are also upgradeable to 1000BaseTx using an upgrade license key. The 2 SFP connectors are always set to 1000BaseTx.

New Power over Ethernet/Power Supply Enhancements

PoE Power Supply Rating Increase

The amount of power available from the PoE power supplies has increased. Note that no hardware upgrade is required for the following new values:

- OS9-IPS-600A increased from 550W to 600W
- OS9-IPS-390A increased from 380W to 390W (also applies to the 6850 High PoE PSU)
- OS9-IPS-230A increased from 230W to 240W (also applies to the 6850 Standard PoE PSU)

OS9-GNI-P24 Power Output Increase

The OS9-GNI-P24 power output has increased from 210W to 260W.

- PoE daughter card is powered in-line (1W per PoE enabled port - **lanpower start**)
- Per port PoE values remain unchanged; 3-16W for the OS6800/OS6850 and 3-18W for the OS9000. The default per port value also remains at 15.4W for all three platforms.

Supported Hardware/Software Combinations

The following table shows the 6.1 software releases that support each of the listed OS9000, OS6850, and OS6800 module types:

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

Module Type	Part Number	6.1.1.R01	6.1.2.R03	6.1.3.R01	6.1.5.R01
OS6800-48L	902378	n/a	supported	supported	supported
OS6800-U24	902351	n/a	supported	supported	supported

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

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

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

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

New Software Features and Enhancements

The following software features and enhancements are new with the 6.1.5.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

Feature/Enhancement Summary

Feature	Platform	Software Package
Increased Number of Authenticated Users	all	base
IPv6 Extensions for BGP	OS6850/OS9000	base
L2 DHCP Snooping Enhancements	all	base
Learned Port Security Enhancements	all	base
RIP Timer Configuration	all	base
Server Load Balancing (SLB) Extended Conditions and Statistics	OS6850/OS9000	base

Software Supported

In addition to the new software features introduced with the 6.1.5.R01 release, the following software features are also supported in 6.1.5.R01, subject to the feature exceptions and problem reports described later in these release notes:

Feature Summary

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

Feature	Platform	Software Package
IPv4 Routing	OS6800/OS6850/OS9000	base
IPv6 Routing	OS6850/OS9000	base
IP DoS Filtering	OS6850/OS9000	base
IPv4 Multicast Switching (IPMS)	all	base
IPv6 Multicast Switching (MLD)	OS6850/OS9000	base
IPv4 Multicast Switching (Proxying)	OS6800/OS6850/OS9000	base
IPv6 Multicast Switching (Proxying)	OS6850/OS9000	base
IP Multinetting	all	base
IP Route Map Redistribution	all	base
IPX Routing	all	base
L2 DHCP Snooping	all	base
L2 Static Multicast Address	all	base
Learned Port Security (LPS)	all	base
Link Aggregation (static & 802.3ad)	all	base
MAC Address Mode	OS9000	base
Multicast Routing	all	base
NTP Client	all	base
OSPFv2	all	base advanced routing
OSPFv3	OS6850/OS9000	base advanced routing
Partitioned Switch Management	all	base
Per-VLAN DHCP Relay	all	base
PIM **PIM-SSM (Source-Specific Multicast)**	all	base advanced routing
Policy Server Management	all	base
Policy Based Routing (Permanent Mode)	OS6850/OS9000	base
Port Mapping	all	base
Port Mirroring (1:24)	OS6800	base
Port Mirroring (1:128)	OS6850/OS9000	base
Port Monitoring	all	base
Power over Ethernet (PoE)	OS6850/OS9000	base
Quality of Service (QoS)	all	base
Redirection Policies (Port and Link Aggregate)	OS6850/OS9000	base
RIPv1/RIPv2	all	base
RIPng	OS6850/OS9000	base

Feature	Platform	Software Package
RMON	all	base
Router Discovery Protocol (RDP)	all	base
Routing Protocol Preference	all	base
Secure Copy (SCP)	all	base
Secure Shell (SSH)	all	base
Server Load Balancing	OS6850/OS9000	base
SSH Public Key Authentication	all	base
sFlow	OS6850/OS9000	base
Smart Continuous Switching **Hot Swap** **Management Module Failover** **Power Monitoring** **Redundancy**	all	base
SNMP	all	base
Source Learning	all	base
Software Rollback	all	base
Spanning Tree	all	base
Syslog to Multiple Hosts	all	base
Switch Logging	all	base
Text File Configuration	all	base
User Definable Loopback Interface	all	base
VLANs	all	base
VLAN Stacking and Translation	OS6850/OS9000	base
VRRPv2	all	base
VRRPv3	OS6850/OS9000	base
Web-Based Management (WebView)	all	base

Feature Descriptions

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 6.1.5 release increases the number of possible 802.1X users to 2K per system, not to exceed 1K per module or stackable unit. This number is a total number of users that applies to all authenticated clients, such as AVLAN and 802.1X supplicants or non-supplicants. In addition, the 6.1.5 release also supports the use of all authentication methods and Learned Port Security (LPS) on the same port.

Access Control Lists (ACLs)

Access Control Lists (ACLs) are Quality of Service (QoS) policies used to control whether or not packets are allowed or denied at the switch or router interface. ACLs are sometimes referred to as filtering lists.

ACLs are distinguished by the kind of traffic they filter. In a QoS policy rule, the type of traffic is specified in the policy condition. The policy action determines whether the traffic is allowed or denied.

In general, the types of ACLs include:

- *Layer 2 ACLs*—for filtering traffic at the MAC layer. Usually uses MAC addresses or MAC groups for filtering.
- *Layer 3/4 ACLs*—for filtering traffic at the network layer. Typically uses IP addresses or IP ports for filtering; note that IPX filtering is not supported.
- *Multicast ACLs*—for filtering IGMP traffic.

Access Control Lists (ACLs) for IPv6

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

source ipv6
destination ipv6
ipv6
nh (next header)
flow-label

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.

Note. 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, AVLAN, and VRRP are *not* discarded.
- **UserPorts**—A port group that identifies its members as user ports to prevent spoofed IP traffic. When a port is configured as a member of this group, packets received on the port are dropped if they contain a source IP network address that does not match the IP subnet for the port. Note that this group is not supported on the OmniSwitch 6800.
- **UserPorts Profile**—In addition to spoofed traffic, it is also possible to configure a global UserPorts profile to specify additional types of traffic, such as BPDU, RIP, OSPF, and/or BGP, to monitor on user ports. The UserPorts profile also determines whether user ports will filter the unwanted traffic or will administratively shutdown when the traffic is received. Note that this profile only applies to those

ports that are designated as members of the UserPorts port group. Note that configuring a UseerPorts profile is not supported on the OmniSwitch 6800.

- **DropServices**—A service group that improves the performance of ACLs that are intended to deny packets destined for specific TCP/UDP ports. This group only applies to ports that are members of the UserPorts group. Using the DropServices group for this function minimizes processing overhead, which otherwise could lead to a DoS condition for other applications trying to use the switch. Note that this group is not supported on the OmniSwitch 6800.

ACL Manager

The Access Control List Manager (ACLMAN) is a function of the Quality of Service (QoS) application that provides an interactive shell for using common industry syntax to create ACLs. Commands entered using the ACLMAN shell are interpreted and converted to Alcatel-Lucent CLI syntax that is used for creating QoS filtering policies.

This implementation of ACLMAN also provides the following features:

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

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.

Authenticated VLANs

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

The 6.1.5 release increases the number of possible AVLAN users to 2K per system, not to exceed 1K per module or stackable unit. This number is a total number of users that applies to all authenticated clients, such as AVLAN and 802.1X supplicants or non-supplicants. In addition, the 6.1.5 release also supports the use of all authentication methods and Learned Port Security (LPS) on the same port.

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

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

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 6.1.5 release provides IPv6 support for BGP using Multiprotocol Extensions. The same procedures used for IPv4 prefixes can be applied for IPv6 prefixes as well and the exchange of IPv4 prefixes will not be affected by this new feature. However, there are some attributes that are specific to IPv4, such as AGGREGATOR, NEXT_HOP and NLRI. Multiprotocol Extensions for BGP also supports backward compatibility for the routers that do not support this feature. This implementation supports Multiprotocol BGP as defined in the following RFCs: 4271, 2439, 3392, 2385, 1997, 4456, 3065, 4273, 4760, and 2545.

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

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.

BPDU Shutdown Ports

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

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

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

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.

DHCP Option-82 (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. See "L2 DHCP Snooping" on page 21 for more information.
- **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.

DNS Client

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

Dynamic VLAN Assignment (Mobility)

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

DVMRP

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

End User Partitioning (EUPM)

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

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).

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.

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.

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 Routing

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
- File Transfer Protocol (FTP)
- Address Resolution Protocol (ARP)
- Internet Control Message Protocol (ICMP)
- RIP I / RIP II
- Static Routes

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

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

Note. 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 Routing

IPv6 (documented in RFC 2460) is designed as a successor to IPv4 and is supported on the OmniSwitch 6850 and 9000. The changes from IPv4 to IPv6 fall primarily into the following categories:

- Address size increased from 32 bits (IPv4) to 128 bits (IPv6)
- Dual Stack IPv4/IPv6
- ICMPv6
- Neighbor Discovery
- Stateless Autoconfiguration
- RIPng
- Static Routes
- Tunneling: Configured and 6-to-4 dynamic tunneling
- Ping, traceroute
- DNS client using Authority records

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

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

IP DoS Filtering

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

- ARP Flood Attack - OS6800/OS6850/OS9000
- Invalid IP Attack - OS6850/OS9000
- Multicast IP and MAC Address Mismatch - OS6850/OS9000
- Ping Overload - OS6850/OS9000
- Packets with loopback source IP address - OS6850/OS9000

IP Multicast Switching (IPMS)

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

Destination hosts signal their intent to receive a specific multicast stream by sending a request to do so to a nearby switch using Internet Group Management Protocol (IGMP). The switch then learns on which ports multicast group subscribers are attached and can intelligently deliver traffic only to the respective ports. This mechanism is often referred to as *IGMP snooping* (or *IGMP gleaning*). Alcatel-Lucent's implementation of IGMP snooping is called IP Multicast Switching (IPMS). IPMS allows OmniSwitch 9000 Series switches to efficiently deliver multicast traffic in hardware at wire speed.

Both IGMP version 3 (IGMPv3), which handles forwarding by source IP address and IP multicast destination, and IGMP version 2 (IGMPv2), which handles forwarding by IP multicast destination address only, are supported. IPMS is supported on IPv4 and IPv6 (MLD) on the OmniSwitch 6850 Series and OmniSwitch 9000 Series. The OmniSwitch 6800 Series only supports IPMS for IPv4.

IP Multicast Switching (IPMS) - Proxying

IP multicast proxying and configuring the IGMP and MLD unsolicited report interval are available with this implementation of IPMS. Proxying enables the aggregation of IGMP and MLD group membership information and the reduction in reporting queriers. The unsolicited report interval refers to the time period in which to proxy any changed IGMP membership state.

IP Multinetting

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

IP Route Map Redistribution

Route map redistribution provides the ability to control which routes from a source protocol are learned and distributed into the network of a destination protocol. A route map consists of one or more user-defined statements that can determine which routes are allowed or denied access to the network. In addition, a route map may also contain statements that modify route parameters before they are redistributed.

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

IPX Routing

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

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

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.

Enhancements to DHCP Snooping provided with the 6.1.5 release 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.

The 6.1.5 release provides the following additional benefits when using the LPS feature:

- A configurable limit to the number of filtered MAC addresses allowed on an LPS port.
- Conversion of dynamically learned MAC addresses to static MAC address entries.
- Support for all authentication methods and LPS on the same switch port.

LPS has the following limitations:

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

Link Aggregation (static & 802.3ad)

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

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

Multicast Routing

The OmniSwitch 9000 switches support multicast routing on IPv4 and includes configuration options for multicast address boundaries, the Distance Vector Multicast Routing Protocol (DVMRP), and Protocol-Independent Multicast (PIM).

Multicast traffic consists of a data stream that originates from a single source and is sent to hosts that have subscribed to that stream. Live video broadcasts, video conferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news services are examples of multicast traffic. Multicast traffic is distinguished from unicast traffic and broadcast traffic.

Multicast boundaries confine scoped multicast addresses to a particular domain. Confining scoped addresses helps to ensure that multicast traffic passed within a multicast domain does not conflict with multicast users outside the domain.

PIM
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.

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*.

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

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.

Only one Port Mirroring session is supported. That session can be configured to a "N-to-1" session where "N" can be a number from 1 to 24 (OS6800) or 1 to 128 (OS6850/OS9000) anywhere on the stack. In other words, you can configure up to 24 or 128 source ports for a single destination port in a session on a stack. You cannot configure port mirroring and port monitoring on the same NI module.

Power over Ethernet (PoE)

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

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

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

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

Quality of Service (QoS)

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

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

- **Basic QoS**—includes traffic prioritization and bandwidth shaping
- **802.1p/ToS/DSCP**—includes policies for marking and mapping
- **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. NAT is not supported.

RIPv1/RIPv2

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

The OmniSwitch 6800/6850/9000 switches support RIP version 1 (RIPv1), RIP version 2 (RIPv2), and RIPv2 that is compatible with RIPv1. In addition, text key and MD5 authentication, on an interface basis, for RIPv2 is also supported.

RIPng

The OmniSwitch 6850/9000 switches support Routing Information Protocol next generation (RIPng) for IPv6 networks. RIPng is based on RIPv1/RIPv2 and is an Interior Gateway Protocol (IGP) best suited for moderate sized networks.

RIP Timer Configuration

The 6.1.5 release provides the ability to configure the following key RIP timer values:

- 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.

Secure Copy (SCP)

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

Secure Shell (SSH)

The Secure Shell feature provides a secure mechanism that allows you to log in to a remote switch, to execute commands on a remote device, and to move files from one device to another. Secure Shell provides secure, encrypted communications even when your transmission is between two untrusted hosts or over an unsecure network.

The OmniSwitch includes both client and server components of the Secure Shell interface and the Secure Shell FTP file transfer protocol. SFTP is a subsystem of the Secure Shell protocol. All Secure Shell FTP data are encrypted through a Secure Shell channel.

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

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

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

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

Secure Shell (SSH) Public Key Authentication

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

Server Load Balancing (SLB)

Server Load Balancing (SLB) software provides a method to logically manage a group of physical servers sharing the same content (known as a *server farm*) as one large virtual server (known as an *SLB cluster*). SLB clusters are identified and accessed at Layer 3 by the use of Virtual IP (VIP) addresses or at Layer 2 or Layer 3 by the use of a QoS policy condition. OmniSwitch 6850/9000 switches operate at wire speed to process client requests addressed to the VIP of an SLB cluster or classified by a QoS policy condition and send them to the physical servers within the cluster.

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

sFlow

sFlow is a network monitoring technology that gives visibility to the activity of the network, by providing network usage information. It provides the data required to effectively control and manage the network usage. sFlow is a sampling technology that meets the requirements for a network traffic monitoring solution.

sFlow is a sampling technology embedded within switches/routers. It provides the ability to monitor the traffic flows. It requires an sFlow agent software process running as part of the switch software and an sFlow collector, which receives and analyses the monitored data. The sFlow collector makes use of SNMP to communicate with an sFlow agent in order to configure sFlow monitoring on the device (switch).

Smart Continuous Switching - OmniSwitch 6800/OmniSwitch 6850

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

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

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

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

Smart Continuous Switching - OmniSwitch 9000

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

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

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

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

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

SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that allows communication between SNMP managers and SNMP agents on an IP network. Network administrators use SNMP to monitor network performance and to solve network problems. SNMP provides an industry standard communications model used by network administrators to manage and monitor their network devices. OmniSwitch 9000 switches support SNMPv1, SNMPv2, and SNMPv3.

Source Learning

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

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

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

MAC Address Mode

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

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

Software Rollback

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

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

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

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

Spanning Tree

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

Q2005 (MSTP) is only available when the flat mode is active for the switch. The flat mode applies a single spanning tree instance across all VLAN port connections on a switch. MSTP allows the configuration of Multiple Spanning Tree Instances (MSTIs) in addition to the CST instance. Each MSTI is mapped to a set of VLANs. As a result, flat mode can now support the forwarding of VLAN traffic over separate data paths.

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

Syslog to Multiple Hosts

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

Switch Logging

The Switch Logging feature is designed to provide a high-level event logging mechanism that can be useful in maintaining and servicing the switch. Switch Logging uses a formatted string mechanism to process log requests from applications. When a log request is received, Switch Logging verifies whether the Severity Level included with the request is less than or equal to the Severity Level stored for the appropriate Application ID. If it is, a log message is generated using the formatting specified by the log request and placed on the Switch Log Queue, and Switch Logging returns control back to the calling application. Otherwise, the request is discarded. The default output device is the log file located in the Flash File System. Other output devices can be configured via Command Line Interface. All log records generated are copied to all configured output devices.

Command Line Interface can be used to display and configure Switch Logging information. Log information can be helpful in resolving configuration or authentication issues, as well as general errors.

Text File Configuration

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

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

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.

VLANs

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

The VLAN management software handles the following VLAN configuration tasks:

- Creating or modifying VLANs.
- Assigning or changing default VLAN port associations (VPAs).
- Enabling or disabling VLAN participation in the current Spanning Tree algorithm.
- Enabling or disabling classification of mobile port traffic by 802.1Q tagged VLAN ID.
- Enabling or disabling VLAN authentication.
- Defining VLAN IPX router interfaces to enable routing of VLAN IPX traffic.
- Enabling or disabling unique MAC address assignments for each router VLAN defined.
- Displaying VLAN configuration information.

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

VLAN Stacking and Translation

VLAN Stacking provides a mechanism for tunneling multiple customer VLANs (CVLAN) through a service provider network over the Ethernet Metropolitan Area Network (EMAN). The service provider network uses one or more service provider VLANs (SVLAN) by appending an 802.1Q double tag or VLAN Translation on a customer port that contains the customer's assigned tunnel ID. This traffic is then encapsulated into the tunnel and transmitted through the service provider network. It is received on another Provider Edge (PE) that has the same tunnel ID. This feature enables service providers to provide their customers with Transparent LAN Services (TLS). This service is multipoint in nature so as to support multiple customer sites or networks distributed over the edges of a service provider network.

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.

Note. Authentication is not supported.

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

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

Web-Based Management (WebView)

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

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

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 6.1.5.R01:

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

No.	Trap Name	Platforms	Description
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.

No.	Trap Name	Platforms	Description
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	—	Two or more slots claim to have the same slot number.
38	alaStackMgrNeighborChangeTrap	—	Indicates whether or not the stack is in loop.
39	alaStackMgrRoleChangeTrap	—	Indicates that a new primary or secondary stack is elected.
40	lpsViolationTrap	all	A Learned Port Security (LPS) violation has occurred.
41	alaDoSTrap	all	Indicates that the sending agent has received a Denial of Service (DoS) attack.
42	gmBindRuleViolation	all	Occurs whenever a binding rule which has been configured gets violated.
43	unused	—	

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

No.	Trap Name	Platforms	Description
59	gmHwVlanRuleTableOverloadAlert	all	An overload trap occurs whenever a new entry to the hardware VLAN rule table gets dropped due to the overload of the table.
60	lnkaggAggUp	all	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.
61	lnkaggAggDown	all	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
62	lnkaggPortJoin	all	This trap is sent when any given port of the link aggregate group goes to the attached state.
63	lnkaggPortLeave	all	This trap is sent when any given port detaches from the link aggregate group.
64	lnkaggPortRemove	all	This trap is sent when any given port of the link aggregate group is removed due to an invalid configuration.
65	pktDrop	all	The pktDrop trap indicates that the sending agent has dropped certain packets (to blocked IP ports, from spoofed addresses, etc.).
66	monitorFileWritten	all	A File Written Trap is sent when the amount of data requested by the user has been written by the port monitoring instance.
67	alaVrrp3TrapProtoError	all	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
68	alaVrrp3TrapNewMaster	all	The SNMP agent has transferred from the backup state to the master state.
69	gmHwMixModeSubnetRuleTableOverloadAlert	—	This trap is not supported in the current release.
70	pethPwrSupplyConflict	—	This trap is not supported in the current release.

Unsupported Software Features

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

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

Unsupported CLI Commands

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

Software Feature	Unsupported CLI Commands
BGP	**ip bgp redist-filter status** **ip bgp redist-filter** **ip bgp redist-filter community** **ip bgp redist-filter local-preference** **ip bgp redist-filter metric** **ip bgp redist-filter effect** **ip bgp redist-filter subnets**
Chassis Mac Server	**mac-range local** **mac-range duplicate-eeprom** **mac-range allocate-local-only** **show mac-range status**
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**
QoS	**qos classify fragments** **qos flow timeout** **show policy classify destination interface type** **show policy classify source interface type**

Software Feature	Unsupported CLI Commands
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**
VLANs	**vlan router mac multiple enable|disable** **vlan binding mac-port-protocol** **vlan binding mac-ip** **vlan binding ip-port**
Chassis Supervision	**show fabric**

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

MIB Name	Unsupported MIB variables
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

MIB Name	Unsupported MIB variables
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

MIB Name	Unsupported MIB variables
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 (RFC 2574)	usmUser
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

Command Line Interface (CLI)

Problem Reports

PR 94134

Display of mobile ports by using the CLI "port range" command shows only the information about the first port in the range on an OS9000.

Workaround: The current CLI output for mobile ports is line based, which does not allow for a practical output of the port range. CLI command to display the mobile port should not use port range in the command.

PR 95896

Changes to user permissions do not have an immediate effect.

Workaround: The Session Manager polls the servers (through AAA) every 5 minutes for changes in permissions. For an immediate impact, the administrator may remove the user from the switch to deprive the user from certain permissions. This way the new permissions are current when the user logs on again. By the same token, whenever an administrator decides to remove a user from the switch, he/she should close all of the sessions to which the user is logged on.

PR 105168

When using the CLI "show ni" command, the display of XFP and SFP Model name and Description displays the manufacturer's name. These fields should display the Model Name and device description.

Workaround: There is no known workaround at this time.

PR 106811

Show interface slot/port "SFP/XFP" field output for a port having SFP plugged in cannot differentiate between 100Fx and Bidirectional SFP & between Gig and CWDM SFP on an OS6850.

Workaround: There is no known workaround at this time.

PR 109841

If filtering is used in command "show ip route", only one gateway will display for ecmp routes.

Workaround: There is no known workaround at this time.

PR 110070

<num> <"string"> <hex> <string> is a standard CLI help option for CLI syntax as input type "string".

Workaround: There is no known workaround at this time.

PR 112720

On OS6800 switches, CLI help options for OS6850 switches are displayed. The OS6850 options are only displayed; they are not considered as valid parameters and will trigger an error message if used.

Workaround: There is no known workaround at this time.

PR 113520

On an OS6800, errors are returned when PoE commands are entered even though PoE is not supported.

Workaround: There is no known workaround at this time

Health Monitoring

PR 110452

The CLI command "interface <number> no l2 statistics" will cause the health monitor to not update the rxtx count.

Workaround: After issuing the "interface <number> no l2 statistics" command, enter the "health statistics reset" command. This will reset all health statistics to zero; the statistics will then begin to update soon thereafter.

RMON

Problem Reports

PR 87683

On an OS6800/OS9000, the RMON object etherStatsPkts65to127Octs and other similar objects from RFC 1757 contain TX and RX packet counts.

Workaround: There is no known workaround at this time.

sFlow

Problem Reports

PR 100009

On an OS6850, sFlowCpTable and sFlowFsTable return data of zero when an existing NI is removed or powered down.

Workaround: There is no known workaround at this time.

PR 106480

The total number of samples generated by B version of the 56504 is less than version A due to ASIC changes.

Workaround: Use a smaller sampling interval for revision B chips.

PR 107462

Sflow Datagrams do not use the EMP port on a OS9000 for IPV4 packets.

Workaround: Use the ports on the slots.

SNMP

Problem Reports

PR 82635

On an OS9000, there is no display of the number or the status of fan modules on an OmniSwitch via SNMP or WebView.

Workaround: The number and status of fans can be displayed via the CLI **show fan** command only.

PR 105290

SNMP and the WebView DVMRP --> Routes page does not display Route Flags.

Workaround: Use the CLI **show ip dvmrp route** command to display Route Flags.

Web-Based Management (WebView)

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.

Problem Reports

PR 85135

In WebView > Health > LED page, the XFP1 and XFP2 LEDS are not displayed.

Workaround: XFP1 and XFP2 LEDs are for realtime packet traffic activity indications. For precise reflection of the Rx/Tx activity, please refer to the corresponding Rx/Tx statistics all through the WebView > Health pages.

PR 89084

For Partition Management, giving a user routing protocols permission, but no IP permission results in some table data not being displayed in the routing pages.

Workaround: Make sure a user gets permission for IP if the user is allowed to view any sort of IP routing protocol pages.

PR 96146

When viewing XOS adjacencies with multiple (2 or more) XoS devices connected to an OS9800 using AOS WebView, will return unknown devices with build 4.4.4.188.B or greater which is correct behavior. With earlier builds, the XOS adjacencies could be returned as an AOS device.

Workaround: There is no known workaround at this time.

PR 96274

On an OS9000, in WebView, Networking > IP Multicast > IPv4 > Switching > Configuration and Networking > IP Multicast > IPv6 > Switching > Configuration pages show the actual configuration only, but not the effective configuration of the system in parenthesis of the Status, Querying, Spoofing, and Zapping.

Workaround: Use CLI to view the effective configuration.

PR 99581

The “scp” command is not supported in WebView.

Workaround: There is no known workaround at this time.

PR 101446

WebView > Physical > Adjacencies home page may show an "Applet Not Found" error when the Java Virtual Machine SDK version is less than 1.2.

Workaround: Upgrade the Java Virtual Machine browser plug-in to a more recent version.

PR 100607

Accessing WebView > Security > Servers using Netscape, a window may appear asking for user name selection on an OS6850 when adding a new server.

Workaround: Close the window and ignore. This window shown is not a WebView pop-up window but a Netscape Form Manager window -- perhaps triggered by the "Server Name" label in the server add windows and being treated as if it was a regular form asking for a user name.

PR 106869

In WebView > Layer 2 > VLAN Mgmt > VLANs Add page, on an OS9000, an error message appears whenever user enters VLAN ranges to be created that include more than 63 VLANs at once.

Workaround: Break desired VLAN range to create into smaller ranges (63 or less). (Note: This configuration setting limit occurs all throughout WebView.)

PR 110482

CLI allows the set up of pollers and samplers in advance of the receiver in order to prevent users from having to set up pollers and samplers configurations every time to advance from one switch to the next; however this is not allowed through SNMP (sFlowMIB) or WebView System > Net Monitoring > sFlow -- Poller (sFlowCpTable) or Sampler (sFlowFsTable).

Workaround: There is no known workaround at this time.

PR 112546

Using the Add button with an existing end user profile slot/port access will overwrite the existing slot/port access.

Workaround: Use the Modify button to add ports to an existing slot.

PR 112908

After upgrading to Java plugin version 1.6 in Microsoft Internet Explorer on Windows platform with configurations for (1) browser to bypass a proxy and (2) for Java Network Proxy Settings to use browser settings or use proxy server, applet may still not show (or proxy login/password box might show up).

Workaround: There are two workarounds:

1) Disable LAN automatic configuration.
   A) open an Internet Explorer browser window.
   B) select from the menu "Tools" > "Options"
   C) select the "Connections" tab

D) click on the "LAN Settings..." button
E) unselect the "Automatically detect settings"
F) click on the "OK" button to close LAN Settings
G) click on the "OK" button to close Internet Options

2) Use direct connection instead (not recommendable for security reasons). On the Windows platform computer machine,
A) go to the Java Control Panel (Start > Settings > Control Panel, click on the Java icon).
B) on the middle "Network Settings" boxed section, click on the "Network Settings" button.
C) select "Direct connection" radio button (the last option)
D) click OK to close Network Settings dialog.
E) click OK to close Java Control Panel.

PR 113092

In WebView, Security > ASA > End User Configuration > Slot/Port Access > Add page, the Slot drop-down is not sorted numerically but as string.

Workaround: Search by slot number as string instead.

PR 113285

In WebView when two or more Internet Explorer browser windows are opened for different routers and any Add, Modify, or Help windows remain opened for one of the routers (say router A), another Add, Modify, or Help window from another router (say router B) might log out the previous router (in this case, router B's newly opened window will log out router A).

Workaround: Two workarounds: 1) Keep track of Open Add, Modify, or Help windows opened; make sure that they are closed before opening another one in any other router browser window(s). 2) Use only one browser window to configure a router at a time.

PR 113515

Using Internet Explorer in WebView, Security > ASA > End User Configuration > Slot/Port Access Add page may not behave as expected especially with the All/None check-box.

Workaround: Use either the corresponding CLI command through a Telnet session or use a Firefox browser instead.

PR 113518

In WebView Internet Explorer browser, while loading IPv6 page, the drop down menu are not ready to be used. The functionality works while the web page is fully loaded. It is fully loaded while the second "Done" message shown in the lower left status bar of the Internet Explorer browser.

Workaround: There is no known workaround at this time.

PR 113566

In WebView Networking > UDP Relay > Services > Destination, if a service does not have a specific destination configured, the service is not shown even though the CLI command “show ip udp relay destination" will list it (there are no functional differences).

Workaround: Refer to either Services > Configuration or Statistics (BootP / Generic Services) to see all current services.

PR 114161

In WebView System > System Mgmt > Install page, the Upgrade C20L section is missing the "Update" button, which prevents a C20L upgrade through WebView.

Workaround: Use the corresponding 'upgrade ni <num> license-key <"string">' CLI command.

PR 114417

WebView help for 'physical >chassis >CMM >hardware component' has broken hyperlinks for Flash Manufacturer and Flash Size.

Workaround: There is no known workaround at this time.

Layer 2

Autonegotiation

Problem Reports

PR 86826

On an OS6800/OS9000, when autonegotiation is disabled and speed is forced to 10 Mbps, copper ports may confuse with a forced 100Mbps link partner and can detect a false link UP. Traffic may not pass in this condition.

Workaround: Enable autonegotiation for this copper port and configure desired speed and duplex settings.

Bridging

Problem Reports

PR 86261

On an OS6800, Ethernet SNAP packets with non-zero Organizationally Unique Identifiers (OUI) will not be classified using port-protocol rule.

Workaround: If applicable, other rules should be used to classify such packets.

PR 94269

The OmniSwitch 9000 counts all error packets as unicast packets in the packets received and error counters, regardless of whether the packet is a unicast packet or a multicast packet. An oversized packet is defined as a packet longer than 9216 bytes. This causes the following behavior in the switch:

1) Received packets longer than 9216 bytes are counted as unicast packets AND as error packets even if the packet is a broadcast or multicast packet.

2) For ports operating at speeds of 10 Mbps or 100 Mbps, a packet is not counted as an error packet unless it is longer than 9216 bytes.

Workaround: There is no known workaround at this time.

PR 94866

Invalid Packets with the DA set to all 0's continue to get bridged by the system.

Workaround: There is no known workaround at this time.

PR 99899

On an OS6800/OS6850, binding rules that require a port number to be specified as part of the rule can no longer be deleted if the port is not set to be mobile.

Workaround: Make the port mobile before deleting the rule.

Ethernet Interfaces

Problem Reports

PR 93114

Layer 2 oversized ingress packets are not eligible for broadcast or multicast. As a result, such packets are treated as unicast packets.

Workaround: There is no known workaround at this time.

PR 95125

On an OS9000, the throughput on a 1 Gbps or 10 Gbps port might be 99.998%. This is due to a variance in the oscillator; the clocking can differ by a few PPM. This variance is accepted by the IEEE standard.

Workaround: There is no known workaround at this time.

PR 95526

On an OS9000, if the admin status of the physical interface is toggled while sending high rates of traffic on the 10G interface, some CRC errors are reported on the host connected to the interface.

Workaround: There is no known workaround at this time.

PR 105079

On OS6850 combo copper ports, the link will toggle once when the SFP is plugged in the corresponding fiber cage.

Workaround: There is no known workaround at this time.

PR 105080

If a Gig Copper SFP is plugged into a combo port on an OS6850, then those combo ports can be used in forced mode only. To use a copper combo port, then set the mode to Forced-Copper. To use the Gig copper SFP in the corresponding copper port, then set the mode of the port to Forced-Fiber.

Workaround: There is no known workaround at this time.

PR 105847

On OS6850-U24x board combo ports with default media preference (i.e. Preferred-Fiber) and with 100M Fiber SFP (Avago HFBR-57E0PZ), the link on fiber combo ports doesn't come up if copper cable is plugged in the corresponding copper ports and the copper link is UP.

Workaround: On OS6850-U24x board combo ports, 100M Fiber SFP (Avago HFBR-57E0PZ) can get a link with fiber cable in Forced-Fiber mode only.

PR 112493

The GNI-C20L fiber Ethernet ports were tested for conformance with IEEE Standard 802.3, 2005 Edition. The test cases use special test equipment to check the behavior of auto-negotiation under a variety of normal and error conditions.

The initial results showed that the port stopped sending the "break link" signal earlier than expected. However, the GNI-C20L completed autonegotiation and established a valid link. These results are currently under investigation. Note that the problem occurred only in this test case and did not occur in normal operation during Alcatel-Lucent product or system testing.

Workaround: This problem should not happen in normal operation in customer environments when auto-negotiation it turned on. Using autonegotiation on fiber ports is recommended.

PR 112958

The GNI-C20L fiber Ethernet ports were tested for conformance with IEEE Standard 802.3, 2005 Edition. The test cases used special equipment to create single bit errors and observe the responses.

The standard allows certain single bit errors to occur in the "idle" sequence that is sent between Ethernet frames. One of the test cases is intended to create the allowable single bit errors in the idle sequence. In this test, the GNI-C20L fiber port incorrectly discarded an Ethernet frame when preceded by an idle sequence containing one of the allowable single bit errors.

Workaround: This test failure does not affect normal operation in customer environments. Single bit errors of the allowable type are very rare in production networks. No invalid frames were accepted. The valid frame that was discarded would be retransmitted by higher level protocols.

PR 114172

Although the CLI allows the user to change the Crossover setting for Fiber Ports on C20L modules, the setting always remains MDI and the port remains operational.

Workaround: Revert the setting back to MDI to avoid any confusion the show command might create.

Group Mobility

Problem Reports

PR 98417

When a MAC is learned on an OS6850 as "Filtered" for one port due to a Group Mobility rule violation, and if the MAC reappears on another port, it will not be updated. The MAC will not be shown as filtered for the new port, but will continue to show filtered in the old port.

Workaround: There is no known workaround at this time.

PR 98661

When using MAC based group mobility rules for VLAN classification on an OS6800/OS9000, some entries are not inserted into the MAC VLAN table. This results in the frame not being classified.

Workaround: The MAC address is not inserted into the table due to MAC address collision in the hashing algorithm. The only workaround is to use another type of mobility rule (non-MAC based) for the VLAN classification for addresses that have collisions.

IP Multicast Switching (IPMS)

Problem Reports

PR 98869

On an OS6800, egressing IP Multicast traffic onto a mobile VLAN from a 6800's mobile port will not remove the 802.1Q frame.

Workaround: There is no known workaround at this time.

PR 105039

When Proxying in IPMS is enabled on an OS9000, IGMPv3 and MLDv2 reports generated by the system on behalf of clients are not aggregated; rather IGMPv3 and MLDv2 reports are generated containing a single record.

Workaround: There is no known workaround at this time.

Link Aggregation

Problem Reports

PR 100537

If the link aggregation state is changed on an OS6850, the config sync status shows incorrectly as synchronised. This will result in the NI's going down on takeover.

Workaround: Issue the **copy working certified flash-synchro** command and then do a takeover.

Port Mirroring/Monitoring

Problem Reports

PR 86338

An OS6800/OS9000 preserves the INGRESS tag format of the packet for EGRESS mirroring, which makes the mirror packet go out tagged though the real egress packet is not tagged.

Workaround: There is no known workaround at this time.

PR 114259

In the overwrite off mode, port monitoring will not put any frames in the pmonitor.enc file. This issue does not occur in the overwrite on mode.

Workaround: There is no known workaround at this time.

Source Learning

Problem Reports

PR 83087

The MAC aging time can take up to twice the configured value to age out a MAC address.

Workaround: There is no known workaround at this time.

PR 94127

In the hardware learning mode, the source MACs from control packets (BPDUs) are learned.

Workaround: There is no known workaround at this time.

PR 94180

In the default learning mode, a MAC address will still be learned as a “bridging” entry instead of the “filtering” entry when it matches the QoS Drop rules. However, the actual packet is being discarded so the operation is functional.

Workaround: The switch could be configured to use “software” learning mode.

PR 94181

If a QoS rule is set on a port to drop all traffic and if LPS is configured on the same port, no MACs are learned as part of LPS, as all the traffic is dropped.

Workaround: There is no known workaround at this time.

PR 94515

On the port which has learned port security enabled, the first packet is not forwarded. Once this packet is validated, subsequent packets are forwarded without any problem.

Workaround: There is no known workaround at this time.

PR 94646

On an OS9000 and OS6850, the MAC address of an untagged packet does not get learned in the default VLAN with "filtering" and with "accept only" tagged frames.

Workaround: There is no known workaround at this time.

PR 95322

A duplicate MAC address is not learned as "filtered" if the entry has already been configured as a permanent entry on a different port. The side effect is that the packet with this MAC from a different port would still be forwarded or flooded.

Workaround: There is no known workaround at this time.

PR 96264

With dynamic link aggregation, even when the link aggregation is in admin state “disable”, the individual ports of the link aggregation exchange LACP packets. Therefore, the MACs from these packets are learned due to hardware learning.

Workaround: There is no known workaround at this time.

PR 97310

Some MAC addresses get flushed from the CMM when the primary port of a link aggregate moves from one NI to another. The total count on the NI is correct.

Workaround: Set the source learning aging time value to a small value.

PR 97670

When the source learning mode is changed on an OS6850, sometimes the number of MACs learned may not be displayed accurately although all the MACs are actually learned.

Workaround: The display count should catch up after one aging time has elapsed.

PR 98053

On an OS9000 EMP Learned Port Security trap, only the IP address of the EMP port of the switch is provided (if the switch has an EMP port) and not the IP address of the offending entity.

Workaround: There is no known workaround at this time.

PR 100761

Disabling a linkAggr port while traffic is running might leave some MAC addresses learned on the primary member port of the linkAggr.

Workaround: Administratively down all ports that transmit across the linkAggr or link down/up the primary port after disabling the link aggregate.

PR 100932

Removing the last member port of a linkAggr while traffic is running can result in MAC addresses learned on that port in the default VLAN of the linkAggr.

Workaround: Administratively down all ports that transmit across the linkAgg or link down/up the port after removing it from the linkAggr.

PR 105399

OS9000 traffic is flooded instead of unicast bridged when the chassis is operating in the distributed source learning mode and the traffic from slots egressing linkaggregation are on different slots and the ingress and egress traffic is asymmetric.

Workaround: Configure the source MAC as static or use link aggregation on the same slot.

PR 106462

When an OS9000 switch has less than 16K MAC address learned and 2 more slots receive data with MAC addresses simultaneously that lead to more than 16K mac on the switch, the CMM displays all the MAC's learned on each NI.

Workaround: There is no known workaround at this time.

PR 106781

On OS9000, sometimes max 16k macs do not learn across linkagg in chassis-distributed.

Workaround: Do not use chassis-distributed mode with linkagg and many MACs.

PR 109764

The "port-security <port-range> mac <mac-address>" command is deprecated. A static MAC can not be added with port-range option in CLI.

Workaround: Use "port-security <port> mac <mac-address>" command. A static MAC can be added to per port configuration.

PR 113559

If we are adding AAA users to the system at a rate that is less than the MAC aging time (for example, endless stress test), then all MACs in the system will not be aged out until the adding is stopped.

Workaround: There is no known workaround at this time.

PR 113671

If a MAC is learned as FILTERED and then later this same MAC comes in on different port, this MAC will not be learned.

Workaround: Wait unit the MAC is aged out, or manually remove this MAC from the mac-address-table. Then this MAC will be learned on the new port.

Spanning Tree

Problem Reports

PR 89316

A BPDU packet with the Root BridgeID of 0xffff...is sent out with every link-up to elicit a BPDU reply from the adjacent switch in the current Auto-Edge Detection mechanism.

Workaround: There is no known workaround at this time.

PR 90297

On an OS6800/OS9000, CST Root convergence in 802.1s may be slow due to the circulation of old ‘good’ spanning tree vectors in the network when a root switch is powered off.

Workaround: 1) Use single MSTP region as much as possible. 2) Tune the performance parameters maxAge and hop count to optimal values for the network.

PR 95308

Temporary traffic loops may occur 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 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.

Workaround: For items 1 and 2 above, there is no known workaround at this time. For item 3, the following workarounds 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 a higher priority (a lower numeric value) for the bridge instead of assigning a lower priority (a higher numeric value) for the existing root bridge.

PR 96358

On an OS9000, during a Spanning Tree reconvergence triggered by link up/down or board power up/down, the DVMRP subsystem may print the following error message: "dvmrpRecvProbe, Looping back our probes".

This happens only if the Spanning Tree protocol selected is RSTP and is caused by rapid transitions of port states. It has been verified that the problem does not happen for switched VLAN traffic. So there is no chance of a real loop in the network. Packets handled in software might experience the problem due to larger transit delays, but does not cause any malfunction.

Workaround: There is no known workaround at this time.

PR 100356

During boot-up, a Link aggregate port can momentarily become forwarding before blocking, causing BPDU and dynamic MACs to be learned on that port. The dynamic MACs could potentially remain learned on the Link aggregate port for some time, resulting in traffic disruption.

Workaround: Manually flush the MACs learned on the blocked port or admin down/up the blocked port.

PR 101214

Toggling the edge port of an OS9000 the very first time after boot-up may cause a TCN to occur in the STP (1d) network.

Workaround: Configure the port as an edge port instead of an autoEdge port or configure the switch to run in RSTP (1w) protocol.

PR 105493

Enabling Spanning Tree on an OS9000 in flat mode after disabling does not work when VLAN 1 is disabled.

Workaround: Enable VLAN 1, enable STP and disable VLAN 1.

Step1: disable Spanning Tree -> vlan 1 stp disable

PR 105788

Some STP entries in the 802.1D standard mib may return out of range or undefined values on an OS9000. This is because we are returning the true values of a newer version of STP (802.1Q 2005), not 802.1D 1998 that the MIB is based on. When a new standard MIB is defined, we can obsolete the old version.

Workaround: There is no known workaround at this time.

PR 106513

On an OS9000, STP can display the wrong port as the next best port. The next best port is not actually used in topology calculation and has no effect on the network.

Workaround: There is no known workaround at this time.

PR 112567

Topology Change Counter and Age values only get updated when there's a port transitioning from blocking to forwarding in the local switch.

Workaround: There is no known workaround at this time. This is how the current implementation works to report the TC Counter and Age.

PR 112571

In an OS6800/OS6850 standalone setup with a large number of VPAs (>3K), there's a chance that the STP topology might not converge at all when a lot of VLANS are applied at the same time. This is due to the fact that the STP NI task doesn't have the CPU resource to receive all BPDUs, causing constant loss of BPDUs, aging, flushing and temporary loops, etc.

Workaround: In this setup, don't apply all 128 VLANS with 28 ports at the same time. Instead, apply 64 VLANS first, wait for 64 STP instances to stabilize, then apply the next 32 VLANS and wait for them to stabilize and then apply the next 32 VLANS.

VLAN Stacking

Problem Reports

PR 102958

On an OS9000, setting SVLAN priority using "vlan svlan priority" command may not work properly.

Workaround: Use the QoS command to trust all VLAN Stacking ports, and use QoS policy to configure SVLAN priority.

Layer 3

Basic IP Routing

Problem Reports

PR 94621

In most cases, routed packets needing fragmentation due to a smaller IP MTU on the egress network will not be fragmented and will be forwarded as-is.

Workaround: There is no known workaround at this time.

DHCP Snooping

Problem Reports

PR 100435

On an OS6850, a mobile port with incoming DHCP traffic is able to have a DHCP binding created for it but it is not shown in the Binding Port Table.

Workaround: There is no known workaround at this time.

PR 106977

When DHCP Snooping is enabled, binding entry is not created against the new root port when STP topology is changed.

Workaround: There is no known workaround at this time.

PR 107186

For DHCP Snooping feature, the CLI should not allow user to configure ip-source-filtering on a trusted port. However, if the port is a member of a link aggregate, the CLI does not display an error and allows the configuration.

Workaround: Make sure the port is not a member of a link aggregate before configuring the ip-source-filtering.

PR 112579

Occasionally some entries from the command "show ip helper dhcp-snooping binding" will display an incorrect lease time. But it is just a display issue, the actual lease time of the entry is valid.

Workaround: There is no known workaround at this time.

PR 113827

If DHCP Snooping is enabled on a DHCP server VLAN, the feature will not work on a server that does not support Option-82.

Workaround: If DHCP Snooping must be enabled on the server VLAN, use a server that supports Option-82, such as a Linux server.

IPv6

Problem Reports

PR 86669

On an OS6850/OS9000, IPv6 Router Advertisement decrementing timers are not supported for prefix valid lifetimes or prefix preferred lifetimes.

Workaround: Use IPv6 Router Advertisement fixed timers.

PR 94546

An OS9000 does not make use of the MTU interface configured via the **IPv6 interface** command and only supports the port MTU. Therefore, even if the configured IPv6 interface MTU is smaller, packets will still be forwarded instead of being dropped and “a packet too big” message returned to the sender. This can cause problems with path MTU discovery.

Workaround: There is no known workaround at this time.

PR 96061

On an OS6850/OS9000, the IPv6 implementation does not follow ICMPv6 RFC 2463 section 2.2 Message Source Address Determination in regard to anycast addresses. Instead IPv6 uses any unicast address configured on the switch.

Workaround: A node that sends an ICMPv6 message has to determine both the source and destination IPv6 addresses in the IPv6 header before calculating the checksum. If the node has more than one unicast address, it must choose the source address of the message as follows:
"If the message is a response to a message sent to a multicast or anycast group in which the node is a member, the source address of the reply must be a unicast address belonging to the interface on which the multicast or anycast packet was received.

PR 99374

On an OS6850, the SNMP alaIPv6NeighborState MIB object has been deprecated. It is replaced by alaIPv6NeighborReachability.

Workaround: All SNMP management previously done using the alaIPv6NeighborState object should switch to the new alaIPv6NeighborReachability object.

PR 105893

A ping6 initiated on an OS9000 to one of its own addresses will succeed, even if the interface on which the address is configured is disabled or inactive.

Workaround: There is no known workaround at this time.

PR 113730

When a Linux client moves from one switch IPv6 interface(i1) to another(i2), it retains the prefixes on i2 for a long time. As a result, the Linux client is not able to access nodes on interface i1. (On fedora-2 kernel 2.6.5-1.358, the prefixes won't flush out even if the cable is unplugged.)

Workaround: (1) Bring down the ethernet port on the Linux client and bring it up to flush out old config. (need to do **ifdown** and **ifup**.); or, (2) Wait until the old prefixes time out.

Server Load Balancing (SLB)

Problem Reports

PR 105700

For SLB clusters on an OS6850 or OS9000, there are no statistics or flow distribution metrics on a server basis. The only server statistic is the number of packets passed to a cluster.

Workaround: There is no known workaround at this time.

Advanced Routing

BGP4

Problem Reports

PR 103524

If there are more than 185 BGP redistribution filters configured in a boot.cfg file from an earlier release on an OS6850 or OS9000, they are not consistently translated into IPRM route-maps that handle the redistribution.

Workaround: Do not attempt to restore more than 185 redistribution filters from a boot.cfg file generated from an earlier release.

DVMRP

Problem Reports

PR 100130

On an OS6800, when interfaces from two different routers that are physically attached to the same network are changed from native DVMRP interfaces to DVMRP tunnel interface on-the-fly or vice versa, the multicast traffic will stop being routed across this path.

Workaround: Disable DVMRP before making the change, and then re-enable it again.

OSPF

Problem Reports

PR 106790

On OS9000, the forwarding address of AS-External LSA (corresponding to gateway of redistributed route) is not updated when there is a change in the OSPF route to the forwarding address, due to deleting OSPF interface configuration.

Workaround: 1) Disable and Enable OSPF admin status after deleting OSPF interface configuration, if redistributed route's gateway was reachable on that OSPF subnet. 2) Disable and Enable IPRM route-map that redistributes routes into OSPF, after deletion of OSPF interface configuration.

OSPFv3

Problem Reports

PR 104078

The route table in OSPFv3 cannot be retrieved via SNMP on an OS9000. There is no support for retrieving the OSPFv3 route table in the official IETF draft MIB for OSPFv3. Displaying the OSPFv3 border router table is also not supported in the IETF draft MIB. This affects all hardware platforms that support OSPFv3.

Workaround: The CLI can display the routing table (show ipv6 ospf routes) and border router table (show ipv6 ospf border-routers) for OSPFv3.

PR 105491

Existing inter-area-prefix and inter-area-router LSAs are not originated into areas that are created after initial convergence of the network on an OS9000. For example, the router is running and it originates a set of inter-area-prefix or inter-area-router LSAs into all known areas. Later, another area is added, the new area will not have any of the inter-area-prefix or inter-area-router LSAs that exist in the other areas prior to the creation of the new area.

Workaround: When adding an area to OSPFv3 after it has been running for a while, it should be globally disabled and then re-enabled using the "ipv6 ospf status disable/enable" command.

PR 105770

Point-to-point interfaces are not supported in OSPFv3 at this time; therefore, 6to4 tunnels cannot be used with OSPFv3.

Workaround: There is no known workaround at this time.

PIM

Problem Reports

PR 102501

Currently, there is no support for an "ip unload" command. This can be a problem for customers who have loaded a DRC loadable module and executed a "write memory" command to update the boot.cfg file. Then decided that they wanted the module permanently removed.

Workaround: To permanently disable a module, execute the module's "disable" command followed by another "write memory" command and a reboot. To also free up any memory possibly used by the disabled module, edit the boot.cfg file using the "vi" command to delete the "ip load <module>" command along with any non-default configuration lines associated with the module.

PR 106343

High multicast traffic rate on an OS6850 running PIM-SM where the RP is on a remote switch may cause high CPU utilization. This is due to the packets being PIM register-encapsulated which are software routed. The high data rate is causing the Register Stop packets to be dropped. Once the Register Stop packets are received, the register-encapsulation will stop and native forwarding will take over.

Workaround: Avoid register-encapsulation by configuring the RP to be on the same switch as the source.

Quality of Service (includes ACLs)

General

Problem Reports

PR 105380

On an OS6800 if there is no ARP entry for the destination, routed traffic matching a drop policy gets dropped in software by the CPU instead of being dropped in hardware.

Workaround: There is no known workaround at this time.

PR 105496

Rules that match multicast traffic do not get logged properly on an OS9000.

Workaround: There is no known workaround at this time.

PR 105764

Due to the method QoS handles condition groups (network, services, etc.), the flushing of conditions corrupts the linkages between the condition and the groups. The condition remain attached to the SLB cluster. However, the group pointers in the condition are invalid. Because of this, the group parameters are restricted.

Workaround: There is no known workaround at this time.

Policy Manager

Problem Reports

PR 94083

Policies, which specify a destination slot/port or destination port group will not be applied to traffic which is routed by the switch, these policies will match only bridged traffic.

Workaround: There is no known workaround at this time.

PR 94125

The OS9000 and OS6850 does not support QoS or ACL rules containing destination port, destination VLAN, or destination MAC on traffic that is routed by the OS9000 or OS6850.

Workaround: There is no known workaround at this time.

PR 95249

On an OS9000, the 802.1p value of IP packets is set to 0. On trusted ports, the 802.1p is not altered for non-ip packets. For IP packets, the prioritization is per the TOS value. The 802.1p is also restamped to reflect the ingress TOS precedence value. If the TOS value is 0, the 802.1p is set to 0.

Workaround: Use QoS 802.1p stamping policies, which match on ingress the 802.1p value, to retain the ingress 802.1p. For example:

```
policy condition c 802.1p 1
policy action a 802.1p 1
policy rule r condition c action a
qos apply
```

PR 99336

On an OS6800, QoS policies which specify ethertype do not match SNAP or 802.3 Raw packets. QoS policies only support Ethernet-II packet format.

Workaround: The policy will match traffic which matches the policy criteria even if the packet is not Ethernet-II if the policy specifies only one of the following: source slot/port, destination slot/port, source MAC or destination MAC.

PR 99931

When tagging a link aggregate on an OS6850, QoS does not trust the individual ports of the link aggregate.

Workaround: Manually set the trust bit of the underlying ports through QoS (qos port <slot/port> trusted), or set the port default to trusted (qos trust ports).

PR 99983

The OS6850 switch cannot boot up properly with a **boot.cfg** that exceeds the QoS limitation. It is not recommended to manually edit the **boot.cfg** to configure your QoS. Booting up with a **boot.cfg** obtained from a "write memory" is fine.

The hardware allocation checking is not done during boot up causing QoS configurations to be out of sync with the hardware capability.

Workaround: To prevent the boot.cfg from going beyond the QoS limitations on a large QoS configuration, proceed as follows: edit a text file with your desired qos configuration, apply the configuration using "configuration apply text_file", and save the boot.cfg using "write memory".

PR 101223

On an OS6800, if a policy rule specifies the keyword "log" or "log interval", then the policy is rejected.

Workaround: Logging is not supported by the OS6800. The keyword "log" and "log interval" has to be removed from the policy rule definition.

Security

General

Problem Reports

PR 89262

NESSUS reports bogus “Vulnerabilities”. Basically, NESSUS collects all those known attacks/vulnerabilities into their test suites.

For example, NESSUS sends: http://<switch-address>/cgi/bin/guestbook.cgi

WebView/HTTP-Server’s response: Prompts user for the default switch login page (which is the normal operation for our embedded server).

Since our HTTP server replies with some form of an HTTP response, NESSUS mistakenly concludes that the HTTP server is vulnerable to this attack.

Workaround: There is no known workaround at this time.

PR 91681

On an OS9000, the following is noted when running a test called "alya.cgi (Backdoors)" in the NESSUS test suite:

Security Note: Web mirroring - http (80/tcp)

The following CGI have been discovered:

Syntax: cginame (arguments [default value])

/web/content/login.html (userName [] password [] B1 [Login])

Workaround: This is the expected behavior for the login pages for WebView and Web-AVLAN authentication. NESSUS is known to provide those bogus reports.

PR 95642

On an OS9000, the Denial of Service testing tool (NESSUS) generates bogus reports.

Workaround: There is no known workaround at this time.

PR 107176

On an OS6800, the router MAC address may incorrectly appear as the source of a DoS attack.

Workaround: There is no known workaround at this time.

PR 111704

Some debug output from "debug systrace appid aaa level debug1" can cause the telnet session to have garbage output.

Workaround: Reset the terminal to recover from this problem.

802.1x

Problem Reports

PR 98375

On an OS6850, DHCP rules are not being used to classify traffic on a regular group mobility port. For this reason, with a matching DHCP rule, Device Classification policy will not consider matching a DHCP rule as having a matching GM rule when the policy is applied.

Workaround: There is no known workaround at this time.

PR 99658

On an OS6850, not all MAC addresses will be learned when testing with a traffic generator to simulate traffic with incremental MAC addresses.

Workaround: There is no known workaround at this time.

PR 100189

On an OS6850, when the MAC address table is full, source learning will not learn MAC addresses dynamically and the non-supplicant table will show more entries because the tables are not synchronized.

Workaround: There is no known workaround at this time.

PR 100614

On an OS6850, when a device is moved from one 802.1x port to another 802.1x port, the device is not classified according to the device classification policy that applies to the new port but is learned on the default VLAN for the new port.

Workaround: Reconfigure the new port by disabling and enabling 802.1x on the port. May also have to reconfigure the device classification policy for the new port as well.

PR 103324

An OS6850 will not change the IP address automatically even if the supplicant client is running that can automatically do the ipconfigure release and renew when dynamically changing classification policy when an IP net rule is configured. Depending on what traffic is running, some packets may satisfy the IP net rule and the supplicant will be classified according to the IP net rule.

Workaround: User has to be aware that when the IP net rule is configured and when dynamically changing the classification policy that as group mobility as one of the classification option, traffic from supplicant may still have the old IP address on the VLAN that the supplicant was classified before the policy is changed. The IP net rule will cause the client to be learned on the VLAN that it was previously learned on. E.g. supplicant is learned on VLAN x and has an IP address with VLAN x's subnet. There is also an IP net rule for VLAN x's IP to be classified on VLAN x. When user dynamically changes the classification policy, the supplicant may still be learned on vlan x because the PC has traffic coming out with VLAN x's subnet and thus device classification task will classify the supplicant on VLAN x again.

PR 106463

The CLI command “802.1x initialize <slot>/<port>” only applies to the supplicants on the specified port. All the supplicants are forced to authenticate again. Non-supplicants on the same port are not affected; no re-classification for non-supplicants is required when this CLI command is used.

Workaround: There is no known workaround at this time.

PR 112173

If an 802.1x port is connected to the switch by a hub, if one connects this port directly to the same switch into different VLAN instead of connecting through the hub, then the 802.1x user will not age out but the 802.1x user is no longer connected.

Workaround: Re-authenticate after the MAC is learned on the new port (with different VLAN).

PR 112338

It is possible that the supplicant may not transit out of the ABORTING state.

Workaround: Reboot the switch.

Authenticated Switch Access

Problem Reports

PR 91812

On an OS9000, the server information displayed with the **show configuration snapshot aaa** command or saved with the **configuration snapshot aaa <file_name>** command contains hashed (encrypted) password/key information. In order for a file created with the latter command to be used for configuring servers, password/key information needs to be edited. AAA expects this information encrypted only at boot-up time while at run time the information should be in plain text. In this particular case, the servers created with configuration apply command could not be used because password/key information is wrong.

Workaround: Always edit password/key information before applying a snapshot file.

PR 107085

Accounting log for scp-sftp displays user IP address as 0.0.0.0 on an OS6850.

Workaround: There is no known workaround at this time.

Authenticated VLANs

Problem Reports

PR 87642

On an OS6800, the CLI command to specifically disable 802.1x or AVLAN authentication on a port will disable either of the authentication options configured on the port.

Workaround: There is no known workaround at this time.

PR 98369

DHCP is not supported with port-binding AVLANs on OS6800/OS6850. When DHCP packets are used to trigger the port binding rules, none of the rules work.

Workaround: There is no known workaround at this time.

PR 106976

When DHCP Snooping's IP Source Filtering is enabled on the Authenticated VLAN port of an OS6850, the authentication (via Telnet or HTTP) will fail.

Workaround: Cannot enable IP Source Filtering on AVLAN ports, since IP Source Filtering (work as expected) is blocking the IP traffic.

PR 108982

MAC address table entry for authentication client is not flushed even after the client logs off.

Workaround: Delete the MAC address manually on the CMM.

PR 113826

Only about 55 AVLAN clients can be authenticated through the WEB access due to the long idle timeout in TCP connections.

Workaround: Use Telnet access for AVLAN authentication.

Policy Server Management

Problem Reports

PR 104283

TACACS+ authenticated user cannot manage the file system simply by enabling read/read-write-file-management privileges.

Workaround: The user must also enable the read-write-services privileges as well to manage the file system.

PR 107086

User with readwrite-scp-sftp privileges is initially queried for authentication and authorization on an OS6850. After login to the scp/sftp shell, only accounting requests are sent to TACACS+ server (if enabled), commands are not queried for authorization.

Workaround: There is no known workaround at this time.

PR 107543

If Loopback0 is defined by the user and there is not a physical route for that IP subnet, a RADIUS client will not be able to communicate with a RADIUS server. As a result, RADIUS authentication will fail as the server is unreachable.

Workaround: There is no known workaround at this time.

PR 111011

TACACS+ authentication via SSH is not available.

Workaround: Use console or Telnet for accessing TACACS+ server.

System

General

Problem Reports

PR 97213

On an OS9000, in all CPU exceptions, a Trace Buffer dump is offered at the beginning of the exception.

Workaround: There is no known workaround at this time.

PR 99862

When an over stressed test caused 100% CPU usage on an OS6850, the "Unable to send running checksum" message was displayed, and the switch was rebooted.

Workaround: Avoid 100% CPU usage.

PR 100127

By default, an OS6850 switch is configured to run in strict-priority. If over-subscription is done on priorities 6 & 7, it will bring down the switch.

Workaround: Over-subscription on priorities 6 & 7 is not yet supported on the switch. Do not configure the over subscription with priority 6 or 7.

PR 113082

When a “show health” command is used on an NI, the CPU utilization displayed is higher than the value displayed for previous releases.

Workaround: There is no known workaround at this time.

PR 113698

The number of Telnet responses do not match the number of “Are You There” commands

Workaround: There is no known workaround at this time.

PR 113923

The **/flash/switch/wv-cert.pem** file and the /**flash/switch/wv-key.pem** file need to be re-created to show the Alcatel-Lucent branding.

Workaround: You can rename the two **wv.pem** files or delete and re-run the install command.

PR 114592

A crash occurred when a secure copy was attempted from a server not running a SSH daemon.

Workaround: Make sure a SSH daemon is running on the server.

Chassis Supervision

Problem Reports

PR 95320

On an OS9000, the Unix **mv** command does not update a file's time stamp. Therefore, the check sum will not detect the change.

Workaround: Use the **cp** command instead.

PR 96225

On an OS9000, the "+++ i2cReadRemoteCMM: Error writing starting address!" error message is seen on the secondary after the primary crashes and a takeover is in progress. However, the system functions normally and takeover still occurs.

Workaround: There is no known workaround at this time.

PR 96327

Extra display character in the swlog. No effect on the switch.

Workaround: There is no known workaround at this time.

PR 96584

On an OS9000, during a reload, when the fabric LED is turned off and on, there is no effect on the switch.

Workaround: There is no known workaround at this time.

PR 98768

If hardware configuration changes are introduced to an OS6800/6850 stack without first ensuring that the Primary software configuration is certified, it is possible to create an endless synchronization cycle.

Workaround: Before making any hardware configuration changes ensure that the Primary is certified. Run the **copy working certified flash-syncro** command before making hardware configuration changes.

PR 98956

An OS6850 supports only single point of failure. Removing multiple stacking cables will introduce multiple points of failure from which the system may not recover.

Workaround: Do not remove multiple stacking cables at the same time.

PR 100825

On an OS9000, if both CMMs are removed from the chassis, Layer 2 local (local to the NI) traffic with learned MACs is switched.

Workaround: Insert at least one CMM into the switch to reset the NIs.

PR 103625

Pulling all the cable simultaneously on a stack of 8 OS6850's causes problem.

Workaround: When pulling the cable, the system will start its topology change. Pull the cable one by one during the topology change will lose the topology information. Try to pull the cable slowly to avoid this problem.

PR 104874

If the CPU is at 100%, then the operation of “copy working certify” or “certify to working” can take more than 12 minutes.

Workaround: Prevent doing a “copy working certified” or “copy working certified flash-synchro” if high CPU utilization or remove CPU load.

PR 113256

Any running OS6800/OS6850 switch that is either the PRIMARY or SECONDARY element within a stack will reboot after trying to change its role to PASSTHROUGH using the "stack clear slot <num> immediate" command.

Workaround: Only run the "stack clear slot <num> immediate" command on IDLE OS6800/OS6850 switches.

PR 113440

In a stacked OS6850 environment, if all the boot.slot.cfg files are removed and power cycled, the console on the primary may be lost and all the elements may not be fully operational. This problem will also be present if you stack several elements and power them on for the first time from the factory.

Workaround: Reboot the system, as the files are automatically regenerated.

PR 113927

One or both of the switch log files becomes so large that there is no free space on the flash drive.

Workaround: Delete swlog1.log and swlog2.log and then run the “swlog clear” CLI command.

Power Over Ethernet

Problem Reports

PR 99583

OS6850 POE units support either 510 or 360 Watt power supplies. If unlike power supplies are mixed, or if an unsupported power supply, such as a 120-Watt power supply are used, a console message and a trap are generated.

Workaround: There is no known workaround at this time.

PR 106121

Currently the test portion of the OS9000 lanpower start command checks available power before any attempt is made to activate.

Workaround: There is no known workaround at this time.

PR 107287

During "update lanpower all" on a stack of 8 OS6850-P24/48, you might see the following error message:

THU OCT 12 14:06:27 : LANPOWER (108) error message:

+++ Unable to Read S19 Response!

Reset Daughter Module!

Done

THU OCT 12 14:48:59 : LANPOWER (108) error message:

+++ Unable to Read BOOT_SECTION_RESPONSE Response!

+++ General Programming Error!

Workaround: Do a "lanpower start" on the NI which failed the firmware update. Once lanpower failed to start which is expected, you then try to "update lanpower <slot>" again on that specific NI.

PR 112402

The disabling of priority-disconnect can sometimes appear not to work. A guard band exists between max-power, and max-power - 5 watts. If the total power consumed on a P24 falls with this guard band, priority-disconnect will not operate if it is disabled. If the total power consumed does not fall in the guard band, priority disconnect will continue to function, even if disable has been selected.

Workaround: There is no known workaround at this time.

PR 114483

It is not possible to set the power priority for a slot.

Workaround: Set the power priority for a single port as follows:

```
-> lanpower 1 priority critical
WRPmiscLanpowerPriority
        slot 1
       port 0
      pri  1

SET request to AppId 108, SnapId 2
  sessionId: 1
  Table: pethPsePortTable (110600)
   Index: 1 0
    Object: pethPsePortPowerPriority (8)
         0 0 0 1 {1}
===>sending...
```

Redundancy / Hot Swap

Hot Swap Time Limitations for OmniSwitch 9000

- All removals of NI modules must have a 30 second interval before initiating another hot swap activity.
- All insertions of NI modules must have a 3 minute interval before initiating another hot swap activity.
- All hot swaps of CMM modules must have a 10 minute interval before initiating another hot swap, reload or takeover activity.
- All takeovers must have a 10 minute interval before following with another hot swap, reload or takeover activity.Problem Reports

Problem Reports

PR 91287

After takeover, the new primary CMM does not keep the DOS statistics held by the previous primary CMM.

Workaround: There is no known workaround at this time.

PR 95840

The whole chassis will reload if more than half of NIs are hot swapped in and out of the chassis at roughly the same time.

Workaround: If hot swap of NIs is required, a user may have to wait until a previous NI has been booted first, and then hot swap the next NI.

PR 96011

On an OS9000, NIs will not reset on an unsynchronized takeover if those NIs are manually reset between configuration change and takeover.

Workaround: There is no known workaround at this time.

PR 98992

The **copy working certified flash-synchro** command display takes a long time on an OS6800 switch. OS6800 does not support tffs (true flash file system) but OS6850 does. So an OS6800 switch takes much longer whenever there is an operation related to file operation.

Workaround: There is no known workaround at this time.

PR 113291

Manual takeover may fail if all the NIs are not ready.

Workaround: Wait for all NIs to receive all configuration commands from the boot.cfg file first. Use the "show module status" before performing a manual takeover.

PR 113668

Under certain rare conditions involving takeover/failover, a dual CMM chassis may have two primaries and all the NIs are shut down by the primary CMM that is not communicating with the NIs.

Workaround: Reload the entire chassis (reload all).

PR 113865

On very rare occasions during Takeover, a large configuration on an OS6800 or OS6850 can cause the system to temporarily run out of buffers resulting in lost events.

Workaround: Reboot the system to recover from this type of failure.

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	+33-388-55-69-29
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.