Release Notes

OmniSwitch 6400/6850/6850E/6855/9000E

Release 6.4.4.R01

These release notes accompany release 6.4.4.R01 software for the OmniSwitch 6400/6850//6850E/6855/9000E hardware. They provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

Note: The OS9600/OS9700/OS9800 are not supported in Release 6.4.4.R01.

ATENTION: Please refer to the 6.4.4.R01 Prerequisite section for important release specific information prior to upgrading.

Contents

Related Documentation

These release notes should be used in conjunction with the OmniSwitch 6400, 6850, 6850E, 6855, and 9000E. The following are the titles and descriptions of the user manuals that apply to this release.

User manuals can be downloaded at:

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

- **OmniSwitch 6400 Series Getting Started Guide**
  Describes the hardware and software procedures for getting an OmniSwitch 6400 Series switch up and running.
- **OmniSwitch 6850/6850E Series Getting Started Guide**
  Describes the hardware and software procedures for getting an OmniSwitch 6850 Series switch up and running.
- **OmniSwitch 6855 Series Getting Started Guide**
  Describes the hardware and software procedures for getting an OmniSwitch 6855 Series switch up and running.
- **OmniSwitch 9000/9000E Series Getting Started Guide**
  Describes the hardware and software procedures for getting an OmniSwitch 9000E Series switch up and running.
- **OmniSwitch 6400 Series Hardware User Guide**
  Complete technical specifications and procedures for all OmniSwitch 6400 Series chassis, power supplies, and fans.
- **OmniSwitch 6850/6850E Series Hardware User Guide**
  Complete technical specifications and procedures for all OmniSwitch 6850 Series chassis, power supplies, and fans.
- **OmniSwitch 6855 Series Hardware User Guide**
  Complete technical specifications and procedures for all OmniSwitch 6855 Series chassis, power supplies, and fans.
- **OmniSwitch 9000E Series Hardware User Guide**
  Complete technical specifications and procedures for all OmniSwitch 9000E Series chassis, power supplies, and fans.
- **OmniSwitch CLI Reference Guide**
  Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.
- **OmniSwitch AOS Release 6 Network Configuration Guide**
  Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethe rnet and VLAN configuration), Layer 3 information (routing protocols), security options (Authenticated Switch Access (ASA)), Quality of Service (QoS), link aggregation.

- **OmniSwitch AOS Release 6 Switch Management Guide**
  Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).
- **OmniSwitch AOS Release 6 Advanced Routing Configuration Guide**
  Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, OSPF, and OSPFv3.
- **OmniSwitch Transceivers Guide**
  Includes SFP and XFP transceiver specifications and product compatibility information.
- **Upgrade Instructions for 6.4.4.R01**
  Provides instructions for upgrading the OmniSwitch 6400, 6850, 6850E, 6855, and 9000E to 6.4.4.R01.
- **Technical Tips, Field Notices**
  Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

- OmniSwitch 6400 Series Release 6.4.4.R01 requires 256 MB of SDRAM and 128 MB flash memory. This is the standard configuration shipped.
- OmniSwitch 6850 Series Release 6.4.4.R01 requires 256 MB of SDRAM and 64 MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 6850E Series Release 6.4.4.R01 requires 512 MB of SDRAM and 128 MB of flash memory. This is the standard configuration shipped.
- OmniSwitch 6855 Series Release 6.4.4.R01 requires 256 MB of SDRAM and 128 MB flash memory. This is the standard configuration shipped.
- OmniSwitch 9000E Series Release 6.4.4.R01 requires 1GB of SDRAM and 256 MB 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 deterine your SDRAM and flash memory.

UBoot, FPGA, Miniboot, BootROM, Upgrade Requirements

The software versions listed below are the minimum required, except where otherwise noted. Switches running the minimum versions, as listed below, do not require any Uboot, Miniboot, or FPGA upgrades when upgrading to AOS 6.4.4.R01.

Switches not running the minimum version required should upgrade to the latest Uboot, Miniboot, FPGA that is available with the 6.4.4.R01 AOS software available from Service & Support.

Note: Refer to the **6.4.4.R01 Upgrade Instructions** document for step-by-step instructions on updating to Release 6.4.4.R01.

OmniSwitch 9000E

Release	Miniboot.uboot CMM	UBoot CMM	UBoot NI	FPGA CMM
6.4.4.R01	6.4.3.479.R01	6.4.3.479.R01	6.4.3.479.R01	Major Revision: 2 Minor Revision: 25 (displays as 0x19; recommended)

Note: The ‘S01’ build is required for ISSU support.

OmniSwitch 6850

Release	Miniboot.uboot	UBoot	FPGA
6.4.4.R01	6.4.3.479.R01 (Minimum) 6.4.4.213.R01 (recommended for OS6850/OS6850E mixed stack)	6.4.3.479.R01 (Minimum) 6.4.4.213.R01 (recommended for OS6850/OS6850E mixed stack)	No minimum requirement

OmniSwitch 6850E

Release	Miniboot.uboot	UBoot	CPLD
6.4.4.R01	6.4.4.213.R01	6.4.4.213.R01	No minimum requirement

Note: Shipped with required versions, no upgrade required.

OmniSwitch 6855 / OmniSwitch 6855-U24X

Release	Miniboot.uboot	UBoot	FPGA
6.4.4.R01	6.4.3.479.R01	6.4.3.479.R01	No minimum requirement

OmniSwitch 6400

Release	Miniboot	BootROM	FPGA
6.4.4.R01	6.4.3.565.R01	6.4.3.565.R01	OS6400-C24/P24 (v16) OS6400-C48/P48 (v11) OS6400-U24 (v10)

Prerequisites: Upgrading to 6.4.4.R01

OS-6850 Auto-negotiation on Combo Ports

In AOS Release 6.4.4.R01 auto-negotiation configuration must be replicated on both fiber and copper mediums for combo ports regardless of the media connection. For example, to disable auto-negotiation on a copper combo port use the following commands to disable both media types:

```
-> interfaces slot/port hybrid fiber autoneg {disable | enable}
-> interfaces slot/port hybrid copper autoneg {disable | enable}
-> write memory
```

Note: Ensure the change is saved in **boot.cfg** file prior to the upgrade.

Forced/Preferred Mode on Combo Ports

Beginning in AOS Release 6.4.4.R01 Preferred-Fiber is the only mode supported on all combo ports. The following hybrid commands are no longer supported:

```
-> interfaces slot/port hybrid forced-fiber
-> interfaces slot/port hybrid forced-copper
-> interfaces slot/port hybrid preferred-copper
```

During the 6.4.4.R01 upgrade of the OmniSwitch any ports configured with the commands above will be converted to the Preferred-Fiber setting and upon boot-up a boot.cfg.err file will be generated.

HIC Server Parameter Change

The '**secret**' parameter has been changed to '**key**' for the '**aaa hic server-name'** command. Perform the following after upgrading if a HIC server is configured :

```
-> aaa hic server-name <servername> ip-address <ip-address> key <key>
```

During the 6.4..4.R01 upgrade the OmniSwitch will automatically convert the '**secret**' parameter to '**key**'.

LPS Port Security Parameter Change

The '**enable'** parameter has been changed to '**admin-status enable'** for the '**port-security'** command. If the **'port-security *slot/port* enable'** command exists perform the following after upgrading:

```
-> port-security <slot/port> admin-status enable
```

During the 6.4.4.R01 upgrade of the OmniSwitch any ports configured with the old command will cause a boot.cfg.err file will be generated.

Note: If the command exists without the '**enable'** parameter**,** the switch will automatically convert the configuration to the new CLI command.

Mixing OS6850/OS6850E Switches in a Stack

The following must be considered before attempting a mixed stack environment:

- To support a mixed stack of OS6850/OS6850Es the OS6850s **MUST BE UPGRADED** to Release 6.4.4.R01 first.
- **In a mixed OS6850/OS6850E stacked evironment you must first upgrade the existing OS6850 switches before adding the OS6850E to the stack.** Additionally, it's recommended that the OS6850 switches be upgraded to U-Boot/Miniboot version **6.4.4.213.R01** to match the OS6850E switches.
- **In a mixed 6850/6850E stack environment never upgrade the FPGA for all the elements at the same time by using the 'all' parameter of the 'update' command.** This will cause all stack elements to use the FPGA version of the Primary element which is not compatible with both models.
- If an OS6850/OS6850E is inserted into a stack with a mode different than the primary element, the inserted switch will not join the stack and will be put into PASS-THROUGH mode. See OmniSwitch 6850E Stacking Mode for additional information on the OmniSwitch 6850E modes.

New Hardware Supported

The OmniSwitch 6850E Series of switches are wire-rate, low latency, stackable, 10-gigabit ethernet switches available in the following models:

OmniSwitch 6850E-24

The OmniSwitch 6850E-24 is a stackable edge/workgroup switch offering the following:

- 20 non-combo 10/100/1000 RJ-45 ports
- 4 combo ports (10/100/1000 RJ-45 or SFP)
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-24X

The OmniSwitch 6850E-24X is a stackable edge/workgroup switch offering the following:

- 20 non-combo 10/100/1000 RJ-45 ports
- 4 combo ports (10/100/1000 RJ-45 or SFP)
- 2 non-combo 10-Gigabit SFP+ ports
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-48

The OmniSwitch 6850E-48 is a stackable edge/workgroup switch offering the following:

- 44 non-combo 10/100/1000 RJ-45 ports
- 4 combo ports (10/100/1000 RJ-45 or SFP)
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-48X

The OmniSwitch 6850E-48X is a stackable edge/workgroup switch offering the following:

- 46 non-combo 10/100/1000 RJ-45 ports
- 2 combo ports (10/100/1000 RJ-45 or SFP)
- 2 non-combo 10-Gigabit SFP+ ports
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-U24X

The OmniSwitch 6850E-U24X is a stackable edge/workgroup switch offering the following:

- 22 non-combo SFP ports
- 2 combo ports (10/100/1000 RJ-45 or SFP)
- 2 non-combo 10-Gigabit SFP+ ports
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-P24

The OmniSwitch 6850E-P24 is a stackable edge/workgroup PoE switch offering the following:

- 20 non-combo 10/100/1000 RJ-45 802.3at PoE ports
- 4 combo ports (10/100/1000 RJ-45 802.3at PoE or SFP)
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-P24X

The OmniSwitch 6850E-P24X is a stackable edge/workgroup PoE switch offering the following:

- 20 non-combo 10/100/1000 RJ-45 802.3at PoE ports
- 4 combo ports (10/100/1000 RJ-45 802.3at PoE or SFP)
- 2 non-combo 10-Gigabit SFP+ ports
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-P48

The OmniSwitch 6850E-P48X is a stackable edge/workgroup PoE switch offering the following:

- 44 non-combo 10/100/1000 RJ-45 802.3at PoE ports
- 4 combo ports (10/100/1000 RJ-45 802.3at PoE or SFP)
- 10-Gigabit SFP+ expansion module

OmniSwitch 6850E-P48X

The OmniSwitch 6850E-P48X is a stackable edge/workgroup PoE switch offering the following:

- 46 non-combo 10/100/1000 RJ-45 802.3at PoE ports
- 2 combo ports (10/100/1000 RJ-45 802.3at or SFP)
- 2 non-combo 10-Gigabit SFP+ ports
- 10-Gigabit SFP+ expansion module

OS6-XNI-U2

The OS6-XNI-U2 expansion module provides 2 SFP+ ports that plugs into the back of an OmniSwitch 6850E chassis in place of the 2 CX4 stacking connectors.

Fanless Power Supply for OS6855-U24X

80W AC fanless power supply for the OS6855-U24X.

900W Power Supply

Provides system and PoE power for the OS6400, OS6850, and OS6850E PoE models:

- 120W of system power
- Up to 780W of PoE power depending on OmniSwitch model

OS9-GNI-P24E

The OS9-GNI-P24E provides 802.3at PoE capability for the OS9000E:

- 24 10/100/1000 RJ-45 802.3at PoE ports

Supported Hardware/Software Combinations

The following table shows the 6.X software releases that support each of the listed OS6400, OS6850, OS6850E, OS6855 and 9000E module types:

Module Type	**Part No.**	**6.1.3.R01**	**6.1.5.R01**	**6.3.1.R01**	**6.3.2.R01**	**6.3.3.R01**	**6.3.4.R01**	**6.4.1.R01**	**6.4.2.R01**	**6.4.3.R01**	**6.4.4.R01**
OS9700E/9702E-CMM	902668	no	no	no	no	no	no	supported	supported	supported	supported
OS9702E-CMM	902808	no	no	no	no	no	no	supported	supported	supported	supported
OS9702-CHASSIS	902727	no	no	no	no	no	supported	supported	supported	supported	supported
OS9-GNI-C24E	902669	no	no	no	no	no	no	supported	supported	supported	supported
OS9-GNI-U24E	902670	no	no	no	no	no	no	supported	supported	supported	supported
OS9-XNI-U2E	902671	no	no	no	no	no	no	supported	supported	supported	supported
OS9-XNI-U12E	902851	no	no	no	no	no	no	no	no	supported	supported
OS9-GNI-P24E	902927	no	no	no	no	no	no	no	no	no	supported
OS6855-14	902648	no	no	no	supported	no	supported	no	supported	supported	supported
OS6855-24	902664	no	no	no	supported	no	supported	no	supported	supported	supported
OS6855-U10	902647	no	no	no	supported	no	supported	no	supported	supported	supported
OS6855-U24	902555	no	no	no	supported	no	supported	no	supported	supported	supported
OS6855-U24X	902802	no	no	no	no	no	no	no	supported	supported	supported
OS6850-24	902457	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-48	902495	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-24X	902458	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-48X	902462	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P24	902459	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P48	902463	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P24X	902460	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P48X	902464	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-U24X	902418	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-24L	902487	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-48L	902489	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P24L	902488	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850-P48L	902490	supported	supported	supported	no	no	supported	no	supported	supported	supported
OS6850E-24	902936	no	no	no	no	no	no	no	no	no	supported
OS6850E-P24	902934	no	no	no	no	no	no	no	no	no	supported
OS6850E-24X	902937	no	no	no	no	no	no	no	no	no	supported
OS6850E-P24X	902935	no	no	no	no	no	no	no	no	no	supported
OS6850E-48	902938	no	no	no	no	no	no	no	no	no	supported
OS6850E-P48	902932	no	no	no	no	no	no	no	no	no	supported
OS6850E-48X	902939	no	no	no	no	no	no	no	no	no	supported
OS6850E-P48X	902933	no	no	no	no	no	no	no	no	no	supported
OS6850E-U24X	902940	no	no	no	no	no	no	no	no	no	supported
6400-24	902621	no	no	no	no	supported	supported	no	supported	supported	supported
6400-P24	902622	no	no	no	no	supported	supported	no	supported	supported	supported
6400-U24	902623	no	no	no	no	supported	supported	no	supported	supported	supported
6400-U24D	902624	no	no	no	no	supported	supported	no	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 an ‘A’ revision level for NI 1:

```
DC-Core ->> show ni 1

Module in slot 1
  Model Name:                    OS9-GNI-C24E,
  Description:                   10 – 1000 RJ45,
  Part Number:                   902669-90,
  Hardware Revision:             F04,
  Serial Number:                 J21Q0772,
  Manufacture Date:              MAY 03 2008,
  Firmware Version:              ,
  Admin Status:                  POWER ON,
  Operational Status:            UP,
  Power Consumption:             51,
  Power Control Checksum:        0x872,
  CPU Model Type   :             Motorola MPC8540 ADS,
  MAC Address:                   00:d0:95:e6:54:80,
  ASIC - Physical 1:             BCM56620_A1
  FPGA - Physical 1:             0007/00
  UBOOT Version :                6.4.3.479.R01
  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:                    OS9802-CMM,
  Description:                   CMM,
  Part Number:                   902672-90,
  Hardware Revision:             B,
  Serial Number:                 J23Q0128,
  Manufacture Date:              MAY 08 2008,
  Firmware Version:              2,
  Admin Status:                  POWER ON,
  Operational Status:            UP,
  Power Consumption:             40,
  Power Control Checksum:        0x9214,
  CPU Model Type   :             Motorola MPC8541 ADS,
  MAC Address:                   00:d0:95:e0:6c:ac,
```

6.4.4 New Software Features and Enhancements

The following software features and enhancements are new with the 6.4.4.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

6.4.4 New Feature/Enhancement Summary

Feature	Platform	Software Package
Access Guardian		
- Accounting for Non-supplicants	all	secu
- Captive Portal Enhancements	all	secu
- Control Over Access Guardian	all	secu
- Dynamic User Network Profiles	all	secu
- Host Integrity Check (HIC) Redundancy	all	secu
Out of the Box Auto-Configuration with Dynamic Management VLAN	all	base
DHCP Option-82 CVLAN	all	base
Dual-Home Links		
- Dual-Home Link (DHL) – Active-Active	all	base
Ethernet OAM		
- Virtual MEP – UNI Loopback	all	base
- Fault Propogation Enhancement	all	base
Link Monitoring/Diagnostics/Recovery		
- Link Monitoring/Flapping Detection	all	base
- Link Fault Propogation	all	base
- Interface Violation Recovery	all	base
- Time Domain Reflectometry	all	base
Learned Port Security Enhancements	all	base
Link Aggregation		
- Minimum LAG size	all	base
LLDP		
- Rogue Detection	all	base
OmniSwitch 6850E Stack Mode	6850E	base
- In 6850 Mode – Supports same software		

Feature	Platform	Software Package
features as OS6850 - In 6850E Mode – Supports same software features as 6855-U24X (VRF/egress policies)		
Power Over Ethernet		
- Auto Negotiation of PoE Class	6850E/9000E	base
- 802.3at support	6850E/9000E	base
Spanning Tree		
- STP Loop Guard	all	base
VLAN-based Ingress Source Filtering / Dynamic ARP Inspection	all	base
Web Cache Communication Protocol (WCCP)	all	base

6.4.4 - New Feature/Enhancement Descriptions

Access Guardian

Accounting for Non-supplicants

Previous releases only supported accounting for 802.1x authentication. This feature allows the same accounting information to be supported for both MAC and Captive Portal authentication.

Captive Portal Enhancements

The following Captive Portal Enhancements have been added:

- **Custom Proxy Port** – Allows an administrator to define a custom proxy port for users being authenticated via Captive Portal.
- **Inactivity Logout Timer –**. When enabled this feature will flush a user from the Captive Portal user table if there is no activity for a set amount of time. The inactivity timer is equal to the MAC aging timer.
- **Public Certificate Support** – This feature allows the administrator to change the name of the Captive Portal URL to match that of a public certificate on the switch. This allows PKA authentication when using Captive Portal.

Control Over Access Guardian Behavior

This feature provides flexibility at the port-level to determine which Access Guardian process is performed first on a device attempting to log on to the network through an 802.1x-enabled port. This flexibility allows the administrator to first apply MAC authentication to the device, even if the device uses 802.1x EAPOL frames for supplicant authentication. After MAC authentication is done, subsequent 802.1x authentication can be applied to the same device.

Applying MAC authentication first allows the system to check if the MAC address of the supplicant device is on a “black list” and should not be allowed to access the network. If the address checks out OK, the device can undergo 802.1x authentication or be classified as a non-supplicant.

Dynamic User Network Profiles (UNP) Enhancement

Currently, users can only be classified in a UNP based on authentication result (802.1X, Captive Portal, or MAC auth) or based on classification rules (IP or MAC ranges). If no authentication mechanisms are configured the switch has no way of assigning a user to a UNP.

This feature enhances the current protocol between the HIC server and the OmniSwitch by allowing the HIC server to return a UNP. A specific user (i.e. MAC address) would then be placed into this UNP based on the information sent. For example, users can then be classified into UNPs based on Active Directory group memberships, machine specific parameters, etc.

Host Integrity Check (HIC) Redundancy

This feature allows the configuration of a primary and backup HIC server (Cyber Gate Keeper) to provide HIC server redundancy. The mode can be configured to determine what happens to users currently in the HIC authentication process when neither of the HIC servers is reachable:

- **Hold** - Hosts stay in their UNP and in a HIC in progress state and do not have network access.

- **Pass-through -** Hosts stay in their UNP but are removed from the HIC in progress state. Hosts have network access according the policy list set for their UNP.

Out of the Box Auto-Configuration (Zero-Touch Configuration)

The Out-of-the-Box Auto-Configuration capability automates and simplifies the deployment of large network installations eliminating the need for manual configuration of each device. It also ensures that each device is compliant with the centrally controlled device configuration policies and firmware revisions.

Learned Management VLAN using Nearest-Edge Mode

An OmniSwitch running the Auto-Configuration feature is automatically enabled to process LLDP PDUs with the unique Nearest-Edge destination MAC address. In Nearest-Edge mode the Management OmniSwitch will use a unique MAC address when sending LLDP PDUs. The Automatic Remote Configuration feature will look for these unique packets to determine a Management VLAN. It will then create a DHCP client interface on that tagged VLAN. The Nearest-Edge mode is useful when a DHCP client interface needs to be configured on a VLAN other than the default VLAN.

DHCP Option 82 with CVLAN

This feature allows for the global configuration of DHCP option-82 Circuit ID and Remote ID in ASCII format. The OmniSwitch currently supports a configurable ASCII string for Circuit ID with the following fields:

- VLAN – The outer vlan
- User string
- System name
- Interface – the slot/port
- Interface alias – the alias configured for the slot/port
- Base mac-address

This feature adds the capability to add the CVLAN in ASCII format for both the Circuit ID and the Remote ID.

Dual-Home Links

Dual-Home Link (DHL) Active-Active

Dual-Home Link (DHL) Active-Active is a high availability feature that provides fast failover between core and edge switches without using Spanning Tree. To provide this functionality, DHL Active-Active splits a number of VLANs between two active links. The forwarding status of each VLAN is modified by DHL to prevent network loops and maintain connectivity to the core when one of the links fails.

This implementation of DHL Active-Active is provided in addition to the previously released LACP-based DHL Active-Standby solution. Both versions are supported. The DHL Active-Active feature, however, is configurable on regular switch ports and on logical link aggregate ports (linkagg ID) instead of just LACP aggregated ports. In addition, the two DHL links are both active, as opposed to the active and standby mode used with LACP.

Ethernet OAM

Virtual UNI Loopback – Virtual MEP

This feature provides support for the configuration of a virtual or loopback MEP that is not attached to a physical switch interface. This eliminates the need to use a physical port for loopback CCM messages.

Fault Propogation Enhancement

This feature is used to propagate OAM Connectivity Fault Management (CFM) events into the interface that is attached to a MEP. This can be used with a point to point Ethernet service between a local UP MEP and a remote UP MEP to propogate a link down event.

Link Monitoring/ Diagnostics/Recovery

Link Monitoring / Link Flapping Detection

The Link Monitoring feature is used to monitor interface status to minimize the network protocol re-convergence that can occur when an interface becomes unstable. To track the stability of an interface, this feature monitors link errors and link flaps during a configured timeframe. If the number of errors or link flaps exceeds configured thresholds during this time frame, the interface is shut down.

There are no explicit Link Monitoring commands to recover a port from a Link Monitoring shutdown; such ports are subject to the interfaces violation recovery mechanisms configured for the switch. The following capabilites are provided:

- **Wait to Restore Time** – Introduces a delay before the interface becomes operational allowing the network to convergence more gracefully.
- **Interface errors monitoring** - Physical errors such as CRC, Lost frames, Errors frames and Alignment errors are monitored. When excessive errors are detected, the interface will be shutdown.
- **Interface flapping** - When excessive interface flapping is detected, the interface will be shutdown.
- **Permanent shutdown** - When an interface has been shutdown too many times it can be placed in this mode requiring it to be enabled administratively.

Link Fault Propagation

The Link Fault Propagation (LFP) feature provides a mechanism to propagate a local interface failure into another local interface. In many scenarios, a set of ports provide connectivity to the network. If all these ports go down, the connectivity to the network is lost. However, the remote end remains unaware of this loss of connectivity and continues to send traffic that is unable to reach the network. To solve this problem, LFP does the following:

- Monitors a group of interfaces (configured as source ports).
- If all the source ports in the group go down, LFP waits a configured amount of time then shuts down another set of interfaces (configured as destination ports) that are associated with the same group.
- When any one of the source ports comes back up, all of the destination ports are brought back up and network connectivity is restored.

Interface Violation Recovery

The OmniSwitch allows features to shutdown an interface when a violation occurs on that interface. To support this functionality, the following interfaces violation recovery mechanisms are provided:

- Manual recovery of a downed interface using a CLI command.

- An automatic recovery timer that indicates how much time a port remains shut down before the switch automatically brings the port back up
- A maximum number of recovery attempts setting that specifies how many recoveries can occur before a port is permanently shutdown
- A wait-to-restore timer that indicates the amount of time the switch waits to notify features that the port is back up
- An SNMP trap that is generated each time an interface is shutdown by a feature. This can occur even when the interface is already shutdown by another feature. The trap also indicates the reason for the violation.
- An SNMP trap that is generated when a port is recovered. The trap also includes information about how the port was recovered.

Time Domain Reflectometry (TDR)

Time Domain Reflectometry (TDR) is a feature that is used to detect cable faults. This feature is best deployed in networks where service providers and system administrators want to quickly diagnose the state of a cable during outages, before proceeding with further diagnosis.

When a TDR test is initiated, a signal is sent down a cable to determine the distance to a break or other discontinuity in the cable path. The length of time it takes for the signal to reach the break and return is used to estimate the distance to the discontinuity.

TDR is an on-demand, out-of-service test. The test is not automatically triggered; data and protocol traffic is interrupted. Only supported on copper ports.

Learned Port Security Enhancements

The following Learned Port Security (LPS) enhancements have been added:

- LPS now continues to learn filtering MAC addresses after the learning window has expired, but only up to the configured filtering MAC address limit.
- A new type of static MAC address (pseudo-static) is maintained. A pseudo-static MAC address is not user-configured; it is a dynamically learned MAC address that is treated the same as a regular static address (will not age out or be flushed during the learning window time period). However, the pseudo-static MAC is not saved in the running configuration.
- New parameter options for the LPS **port-security shutdown** CLI command.
    1. **No Aging of Learned MAC Addresses**. A new **no-aging** parameter specifies whether or not LPS will learn MAC addresses as “pseudo-static” addresses.
    2. **Convert MAC Addresses to Static MACs**. A new **convert-to-static** parameter specifies whether or not pseudo-static and dynamically learned MAC addresses are converted to static MAC addresses when the learning window time expires.
    3. **Learning Window Start at Boot-up**. A new **boot-up** parameter specifies whether or not LPS will start the learning window time when the switch boots up.
- New **admin-state** parameter for the **port-security** CLI command. This parameter is used to enable, disable, or lock an LPS port. In addition, the **port-security** command will now accept a range of ports.
- Creating a static MAC address on a port now automatically enables LPS on that port.

- New **brief** parameter for the **show port-security** CLI command. This parameter is used to provide a summary of the LPS status, configuration, and MACs learned on all the LPS ports.
- The VLAN ID bound to an LPS static MAC address is automatically updated when the default VLAN for the LPS port is changed.
- Duplicate LPS static MAC addresses are now allowed on different ports within the same VLAN. However, dynamic MAC addresses that match a configured static MAC address within the same VLAN are not learned.
- The "Bridge MAC Learned" and "LPS Violation" SNMP traps now have three fields of information: port number, VLAN ID, and MAC address.
- A new LPS shutdown violation mode, "discard", is now supported. This mode administratively disables the port, but the port remains physically up. The "shutdown" and "restricted" modes are still supported.

Link Aggregation

Minimum Link Aggregation Size

Allows an administrator to configure a minimum number of ports to be active on a link aggregate.

- When number of attached ports is below the minimum size the aggregate is automatically disabled.
- When number of attached ports returns above the minimum size the aggregateis automatically re-enabled.

LLDP Rogue Detection

LLDP rogue detection provides secure access to the network by detecting rogue devices and preventing such devices from connecting through any OmniSwitch port. A trusted LLDP agent can be assigned to individual ports, slots or the whole chassis. A trusted agent can be assigned by configuring the chassis ID sub type that will be used to validate the chassis ID type of the incoming LLDPDU.

The port can be moved to a violation state and a trap and/or port shutdown can be configured when the following instances occur:

- If more than one LLDP remote agent is learned on a port
- If no LLDPDU is received within 3 times the LLDP transmit interval (30 seconds) after link activation on a port that has no trusted remote agent configured
- If the same chassis ID and port ID of the remote agent already exists in the trusted remote agent database but on a different port.

OmniSwitch 6850E Stacking Mode

The OmniSwitch 6850E can be configured to run in either of two modes. Depending on the mode it can operate as a standalone switch or be stacked with other 6850 or 6850E models as shown in the table below:

Mode	Capability	SW Features
6850 (Default)	**- Allows stacking with OS6850 using CX4 module.** - Allows stacking with OS6850E using CX4 or SFP+ module.	Same as OS6850
6850E	- Cannot stack with OS6850 models. **- Allows stacking with OS6850E models using CX4 or SFP+.**	Same as OS6855-U24X (Including VRF and Egress Policies)

Note: To support a mixed stack of OS6850s and OS6850Es, the OS6850s **MUST first be upgraded to AOS Release 6.4.4.R01.** Refer to the 6.4.4.R01 prerequisites section and the Upgrade Instructions for more detailed information.

Note: If an OS6850/OS6850E is inserted into a stack with a mode different than the primary element, the inserted switch will not join the stack and will be put into PASS-THROUGH mode.

Power Over Ethernet

Support for 802.3at, Automatic Power Class Detection, and 900W Power Supply

These features allow the OmniSwitch to provide up to 30W of PoE power as well as automatically detect the Class (Class 0, Class1, Class2,Class3 or Class4) of the connected powered device. This allows the OmniSwitch to automatically adjust the maximum allowed power for a port preventing the OmniSwitch from delievering more power than the device requires. Refer to the table below for PoE feature and power support.

	OS6400	OS6850	OS6855	OS6850E	OS9-GNI-P24E
Maximum PoE Power	780W	480W	80W (C24) 66W (C14)	780W	720W
Maximum PoE per Port	18W	18W	20W	30W	30W
802.3at Support	No	No	No	Yes	Yes
PoE Class Detection	No	No	No	Yes	Yes

Note: Maximum PoE power available for OS6400, OS6850, OS6850E based on 900W power supply.

Spanning Tree

STP Loop Guard

This feature is intended to prevent loops in a spanning tree bridged network when a device is unable to receive BPDUs on a non-designated port in a timely manner.

Loop formation can occur when a bridge hosting a blocking port transitions that port to forwarding erroneously. This can lead to a temporary or even a permanent loop.

This feature can be enabled either on a port or link aggregate and can be configured for any spanning tree mode (flat, 1x1, STP, RSTP, MST, PVST). Loopguard effectively protects each STP instance when configured on a port that supports multiple spanning tree instances.

VLAN Based Ingress Source Filtering (Dynamic ARP Inspection)

The 6.4.4.R01 Release introduces VLAN-based Ingress Source Filtering to provide the ability for the user to configure ISF at the VLAN level. When ISF is enabled at the VLAN level the switch will attempt to match the VLAN of the packet received along with matching IP/MAC/Port. When ISF is enabled at the VLAN level the switch will only accept packets if they match the IP/MAC/PORT combination which is obtained from the DHCP snooping binding table entry. All other non-matching packets will be dropped on the ingress port irrespective of the VLAN.

Web Cache Communication Protocol (WCCP)

WCCP enables the OmniSwitch to transparently redirect traffic to a cluster of cache-servers. The server can be a web cache engine or any kind of cache engine. WCCP allows utilization of web cache engines (or other caches running WCCP) to localize web traffic patterns in the network, enabling content requests to be fulfilled locally. Traffic localization reduces transmission costs and download time. A WCCPv2 enabled switch would redirect the traffic on configured protocol (TCP/UDP) ports on the cache engine instead of the intended hosts directly. WCCP Protocol involves two major functions:

- It allows the WCCP enabled router for transparent redirection to discover, verify, and advertise connectivity to one or more cache servers. This would allow deploying cache servers without the need to reconfigure the cache-server at the client location.
- It allows the designated web-cache to dictate how the router distributes redirected traffic across the cache server cluster.

Previous Release 6.4.3 – Features and Enhancements

The following software features and enhancements were introduced in the 6.4.3.R01 release.

Note: New Default Switch Behavior in 6.4.3 due to Out of the Box Auto-Configuration feature.

Newly deployed or upgraded switches with no *boot.cfg* file running AOS 6.4.3 will automatically run the Out of the Box Auto-Configuration feature. This causes the CMM OK/OK1 LED to blink amber while the process is running. If the Auto-Configuration process is not successful the CMM OK/OK1 LED will continue to blink amber as long as no *boot.cfg* file is on the switch, this is normal behavior in 6.4.3.

Additionally, the Auto-Configuration feature will automatically create a **dhcp-client** IP interface on VLAN 1. This interface can be deleted using the '**no ip interface dhcp-client'** command if desired.

Once the Auto-Configuration process times out (approximately 30 seconds) the switch configuration can be saved to the *boot.cfg* file using the '**write memory'** command. The CMM OK/OK1 LED will then turn solid green as in previous releases.

To ensure the Auto-Configuration process is able run properly so that the **'write memory'** command can be entered , at least one NI MUST be inserted in the chassis-based OS9000 and OS9000E switches. (PR 148181)

6.4.3 Feature/Enhancement Summary

Feature	Platform	Software Package
AAA/802.1x		
- Service Type information in RADIUS Access Request	all	base
- Capture Client IP in RADIUS Accounting Message	all	base
Access Guardian		
- Javaless Captive Portal and MAC OS Support	all	encrypt
Out of the Box Auto-Configuration	all	base
DHCP		
- Internal DHCP Server	all	base
- DHCP Client with configurable option 60	all	base
- DHCP Option 82 ASCII support	all	base
Ethernet OAM		
- IEEE 802.1ag Version 8.1	all	base
- ITU Y.1731	all	base
- Service Assurance Agent (SAA) for OAM and IP SLA Measurements	all	base

Feature	Platform	Software Package
Ethernet Services		
- L2 Control Protocol Tunneling (L2CP)	6400/6850/6855/9000	base
- Wire-Speed Ethernet Loopback	6400/6850/6855/9000	base
- SVLAN Routing	all	base
IP Enhancements		
- Extended Ping & Traceroute	all	base
- Selectable IP Interface for Management Services	all	base
- IP Loopback0 Address In the Same Range of Existing Subnet	all	base
Link Aggregation		
- Non-unicast Load Balancing on Link Aggregation	all	base
- Active-Stand by Dual Home LinkAgg	all	base
LLDP Network Policies	all	base
- Voice Vlan Support	all	base
- Voice Application Support	all	base
MAC-Forced Forwarding (RFC 4562)	all	base
Multiple VLAN Registration Protocol (MVRP)	all	base
Multicast Switching and Routing		
- VRF Aware Multicast Routing (PIM)	6855-U24X/9000E	advanced routing
QoS		
- Egress Policy Rules	6400/6855-U24X/9000E	base
- sr-TCM and tr-TCM (RFC 2697/2698)	all	base
- IEEE 802.1q/ad CFI/DEI Bit Stamping	all	base
- Policy Condition Enhancements (VLAN group, 802.1p Range)	all	base
- Flexible Inner DSCP/ToS Mapping to Outer 802.1p	all	base
- QOS Statistics	all	base
Recursive Static Route	all	base
Security		
- BPDU Shutdown Auto-Recovery Timer	all	base
- Admin User Remote Access Restriction Control	all	base
Storm Control		
- Extended Flood Control Metering for Unknown Unicast, Multicast and Broadcast	all	base

Feature	Platform	Software Package
USB Support	all	base

Previous Release 6.4.3 - Feature/Enhancement Descriptions

AAA RADIUS

Service Type Information in Radius Access Request

The OmniSwitch will add the Service Type attribute in the Access Request to be used by the RADIUS server to distinguish between different request types.

Access Guardian

Captive Portal

- **Javaless – OS Agnostic** – To enhance the number of supported platforms Captive Portal no longer uses Java scripts for releasing or renewing IP addresses.

- **MAC OS Support** - Captive Portal is now supported on MAC OS using Safari version 4. The table below provides all platform and browser support for Captive Portal.

- **Authentication Redirect URLs** - Captive Portal provides the ability to redirect users to different URLs based upon successful or failed authentication.

- **Configurable DNS dictionary –** By default Captive Portal replies only to DNS packets that contain one of the following pre-defined DNS strings: www, http, proxy, wpad, captive-portal, go.microsoft, Mozilla. Starting 643 these keywords can be replaced or augmented by the network administrator.

- **Customizable Banner** – A customizable banner image can be configured that Captive Portal will display at the top of all pages.

Platform	Web Browser Software
Windows 2000, Windows XP, and Windows Vista	Internet Explorer Firefox 3
Linux	Firefox 3
Mac OS X 10.5 Leopard	Safari version 4 Firefox 3

Captive Portal Browser Support

Out of the Box Auto-Configuration (Zero-Touch Configuration)

The Out-of-the-Box Auto-Configuration capability automates and simplifies the deployment of large network installations eliminating the need for manual configuration of each device. It also ensures that each device is compliant with the centrally controlled device configuration policies and firmware revisions.

This feature allows a newly deployed OmniSwitch to automate the process through an instruction file that provides the necessary actions to download its configuration or any necessary firmware upgrades with no user intervention by doing the following:

1. Automatically configures the switch with a DHCP client interface on VLAN 1.
2. Lease an IP address, mask, default gateway, and system name from a reachable DHCP server.

3. Download an instruction file with information to obtain the configuration file, image files and/or script files from given TFTP, FTP or SCP servers.
4. Download and apply the image and configuration file.
5. Automatically reboot with the upgraded image files and switch configuration file or if no images or boot configuration is downloaded scripted instructions are executed on the fly and the switch is made available remotely.

Note: New Default Switch Behavior in 6.4.3 due to Out of the Box Auto-Configuration feature.

Newly deployed or upgraded switches with no *boot.cfg* file running AOS 6.4.3 will automatically run the Out of the Box Auto-Configuration feature. This causes the CMM OK/OK1 LED to blink amber while the process is running. If the Auto-Configuration process is not successful the CMM OK/OK1 LED will continue to blink amber as long as no *boot.cfg* file is on the switch, this is normal behavior in 6.4.3.

Additionally, the Auto-Configuration feature will automatically create a **dhcp-client** IP interface on VLAN 1. This interface can be deleted using the '**no ip interface dhcp-client'** command if desired.

Once the Auto-Configuration process times out (approximately 30 seconds) the switch configuration can be saved to the *boot.cfg* file using the '**write memory'** command. The CMM OK/OK1 LED will then turn solid green as in previous releases.

To ensure the Auto-Configuration process is able run properly so that the **'write memory'** command can be entered , at least one NI MUST be inserted in the chassis-based OS9000 and OS9000E switches. (PR 148181)

DHCP

Internal DHCP Server Functionality

The OmniSwitch now supports an internal DHCP Server compliant with RFC 2131 based on Vital QIP 5.6 release. This feature can be used to provide IP addresses for small offices, management network, or local phone services including support for option 60 and option 43.

Note: For switches shipped with AOS release 6.4.3 the following two templates are pre-loaded on the switch and can be used as examples. If upgrading to 6.4.3 the template files can be downloaded from the Service & Support website:

- *dhcpd.conf.template*
- *dhcpd.pcy.template*

DHCP Client Interface with option 60

The Omni Switch now supports DHCP client functionality on any one configured VLAN. The DHCP client configured interface on an OmniSwitch can obtain an address from a DHCP server and create an IP interface for that VLAN on the switch.

- Release / Renew
- Lease Time
- Automatically configured the learned router as the switch's default gateway.
- Option 60 is configurable and it is sent as part of DHCP discovery/request packet
- Option 12 can be use to configure the switch's system name

DHCP Option 82 ASCII

When the OmniSwitch is configured to stamp DHCP option-82 can be configured to provide a flexible ASCII string for the Circuit-ID value.

Ethernet OAM

The OmniSwitch now supports Ethernet OAM 802.1ag Version 8.1 and ITU Y.1731.

ETH-LB/DMM
ETH-Loopback and ETH-DMM can be used to measure delay and jitter. ETH-DMM can measure by sending out frames with DM information to the peer MEP and receiving frames with DM information from the peer MEP. The ETH-LB test output was improved to look like standard ping providing on demand information for round-trip delay and a summary of min/avg/max delay.

Service Assurance Agents (SAA)
The OmniSwitch's Service Assurance Agents (SAAs) gives users the ability to verify service guarantees, increase network reliability by validating network performance, proactively identify network issues. Service Assurance Agent uses active monitoring to generate traffic between MEPs in a continuous, reliable, and predictable manner, thus enabling the measurement of network performance and health.

The SAA agent is extended to support IP-SLA meassurements using icmp with plans to include udp and tcp support.

Ethernet Services – L2 Control Protocol Tunneling (L2CP)

L2 Control Protocol Tunneling
Enhances the User Network Interface (UNI) profile to allow control packets for the OAM, MVRP and Lacpmarker protocols to be tunneled through the provider network with the configured destination MAC address.

Additionally, support for the following Cisco protocols is added: VTP, VLAN, Uplink Fast, UDLD, PAGP, DTP, CDP.

Wire-speed Ethernet Loopback
A wire-speed Ethernet loopback test function is available to perform In-Service and Out-of-Service throughput testing during initial turn-up or on-the-fly in an active network. The loopback tests can be used to validate the configured Service Level Agreements (SLAs) and QoS parameters that are associated with a service or a flow..

SVLAN Routing
SVLANs now support routing of IPv4 traffic. IPv6 is not supported.

IP Enhancements

Extended Ping & Traceroute Functionality

Ping and Traceroute have been enhanced to allow for additional parameters to be specified.

Ping:
- Set the Source IP
- Set TOS value
- Set DF bit in IP header
- Set data pattern
- Set sweep range

Traceroute:
- Set the Source IP
- Set Timeout in seconds
- Set Probe count
- Set Min and Max TTL
- Set Port number

IP Managed Interfaces

Provides ability to configure a permanent source IP interface to be used when sending packets. The source IP interface can be the Loopback0 address or an existing IP interface on the switch and can be defined for the following applications:

- DNS, FTP, LDAP-SERVER, NTP, RADIUS, SFLOW, SNMP, SSH, SYSLOG, TACACS, TELNET, TFTP

Loopback0 IP in Same Range of Existing IP Interface

The Loopback0 address can now be configured in the same range as an existing IP interface on the OmniSwitch.

Link Aggregation

Non-Unicast Load Balancing on Link Aggregation

The OmniSwitch now supports load balancing of non-unicast (broadcast, multicast, flood) traffic over Link Aggregation. Hashing criteria is configurable.

By default the hashing keys are derived from the flow-based attributes listed below:

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

Active-Standby Dual Home Link

Dual Home Link feature is an edge technology that allows a switch to have redundant connections to two different core/distribution boxes without depending on STP to protect the links providing sub-second convergence times. The edge switch is configured with a link aggregation of size 2 in which one port is configured in standby mode.

The protection is triggered based on detection of the primary link failure and recovery can be controlled and scheduled according to given configuration parameters. It is also possible to stay in the former standby link to avoid additional network outages when primary link recovers. Pre-emption can be enabled to allow primary link to become the active forwarding link upon recovery.

Note: Convergence time varies based on the platform and interface types.

LLDP Network Policies

LLDP Network policy allows the advertisement of VLAN id, 802.1p and DSCP for the following applications: Voice, Voice Signaling, Guest Voice, Guest Voice Signaling, Soft phone voice, Video Conferencing, Streaming voice and Video Signaling.

The OmniSwitch use LLDP-MED Network Policies to advertise the Voice VLAN to the connected IP Phones through explicit definition of LLDP-MED Network Policy that contains information about the VLAN-ID and the associated L2 and L3 priorities. The binding of the network policies can be done globally or on a per port basis. The VLAN must be created explicitly. When using authenticated or mobile VLANs it is recommended to use mobile-tag rules to dynamically associate the devices according to the incoming tagged traffic.

MAC-Forced Forwarding (Dynamic Proxy ARP)

MAC-Forced Forwarding (Dynamic Proxy ARP) is a mechanism to ensure the L2 separation of stations in the same VLAN beyond the local switch. The current port mapping functionality is limited to isolate user ports in the same switch. With MAC-FF the capability is extended to shared topologies such as rings or daisy chains to prevent users from communicating directly and ensuring that all communication happens via their default gateway. In order to accomplish this, the OmniSwitch supports Dynamic Proxy ARP which combines the functionality of port mapping and dhcp-snooping to dynamically learn a router's addresses and act as a local arp proxy for the VLAN's router. Dynamic Proxy ARP - MAC Forced Forwarding uses the following features:

> **Port Mapping** - Port Mapping forwards traffic from user-ports only to network-ports, preventing communication between L2 clients in the same VLAN in the same switch. This prevents direct communication between clients in the same VLAN forcing all traffic to be forwarded to the head end router.
>
> **Dynamic Proxy ARP** - All ARP requests received on port mapping user-ports are answered with the MAC address of the head end router. Dynamic Proxy ARP dynamically learns the IP and MAC address of a head end router and responds with that router's MAC address instead of flooding the ARP request.
>
> **DHCP Snooping** - Snoops the DHCP packets between the server and clients. DHCP snooping is used to dynamically learn the IP address of the head end router.

Multiple VLAN Registration Protocol (MVRP)

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

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

Note: Starting in 6.4.3 MVRP is the default mode for VLAN registration.

QOS

QoS Egress Policy Rules

Omniswitch egress policy rules allow adminitrators to enforce traffic controls on the egress queues as a "last resort" action. By default, QoS policy rules are applied to traffic ingressing the port. The QoS Policy List feature includes an "egress" policy list option to create a list of rules that are applied to traffic egressing a destination port(s). If a policy rule is not associated with an egress policy list, the rule will only apply to ingress traffic.

Tri-Color Marking

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

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

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

There are two types of TCM marking supported:

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

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

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

IEEE 802.1q/ad CFI/DEI Bit Stamping

When sr/trTCM ingress rate limiter is used, frames that are non-conforming to the SLA (yellow) might still be delivered to the egress port when the port is not congested. By enabling CFI/DEI bit stamping on these frames, a color-aware upstream switch would be able to treat these frames differently and drop them first when the network is congested.

QoS Policy Condition Enhancements

- VLAN IDs can be grouped together into a single VLAN group. Similar to other QoS group types, such as MAC and port groups, creating a VLAN group avoids having to configure a separate policy condition for multiple VLAN IDs.
- Specifying a range of 802.1p values for a policy condition is now supported. A range of values is supported when configuring both inner and outer 802.1p policy conditions. A condition must use either a single 802.1p value or a range of 802.1p values; both are not supported at the same time.

Map Several Inner DSCP/ToS Values to the Same Outer 802.1p Value

The ability to specify a range of 802.1p values is particularly useful when classifying Ethernet Services SAP traffic. A new option in a SAP profile suspends the use of SAP bandwidth and priority actions. This allows the use of QoS rules for advanced classification of SAP traffic, such as mapping several DSCP/ToS values to the same outer 802.1p value.

QoS Statistics Enhancements

- QoS statistics monitoring allows the gathering of egress CoS drop and transmit packet statistics for individual ports. Enabling this type of monitoring also allows the user to display egress CoS queue statistics on a per port basis using existing QoS show commands.
- Tri-Color Marking (TCM) policy action now includes a counter color mode option. This option determines which metered packets are counted based on the color the packet was marked by the TCM policy. Enabling this option also allows the display of the counter color statistics using existing QoS show commands.
- QoS commands used to display traffic statistics and system resource usage now include statistics for egress traffic. This applies to traffic classified using egress policy rules.

Recursive Static Route

Recursive static routes are similar to static routes. However, with a recursive static route the gateway does not have to be a directly connectedto the router. If the OmniSwitch is unable to find a route in the routing table for a packet it can use the recursive static route. The OmniSwitch will use its routing table to lookup a route for the gateway instead of having to use a directly connected router. This feature can be used in large networks to configure a uniform static route for all routers on a network. Each router will use the same gateway but the path to reach the gateway may differ for each router.

Security

Admin User

The OmniSwitch can be configured to allow the admin user to only have access to the switch via the console port.

BPDU shutdown auto-recovery timer

Allows ports that are configured in the UserPorts port group to be automatically re-enabled after receiving a spanning tree BPDU.

Storm Control

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

USB Support

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

- Disaster Recovery – The switch can boot from the USB drive if it is unable to load AOS from flash.

  Note: Disaster Recovery requires a minimum 6.4.3 version miniboot/uboot revision to operate.
- Upload / Download Image and Configuration Files - To create or restore backup files.
- Upgrade Code - Upgrade code with the image files stored on the USB drive.

VRF

PIM-DM and PIM-SM are now VRF aware.

Previous Release 6.4.2 - Features and Enhancements

The following software features and enhancements were introduced in the 6.4.2.R01 release.

6.4.2 Feature/Enhancement Summary

Feature	Platform	Software Package
10Km Stacking	OS6855-U24X	base
802.1x Radius-down Fail-Open	all	base
DDM - Transceiver Digital Diagnostic Monitoring	all	base
DHCP Snooping Option 82 – Port-based format	OS6400/OS6850/OS6855	base
ECMP – Support for up to 16 paths	OS6850/OS9000/OS9000E	base
Ethernet Services		
- L2 Tunneling Enhancements	all	base
- Egress Rate Limiting	OS6400/OS6855-U24X/OS9000E	base
Ethernet OAM 802.3ah – EFM	OS6400/OS6850/OS6855	base
Ethernet Ring Protection (ERP) – Shared VLAN	all	base
IGMP Relay - Forward to Specific Host in L3 Environment	OS6850/OS9000/OS9000E	base
IPMVLAN Group Address and Mask	OS6400/OS6850/OS6855	base
MPLS		
- VPLS Support	OS9000E	mpls
- MPLS Static Fast Re-Route	OS9000E	mpls
- MPLS License	OS9000E	mpls
- MPLS OAM-LSP Ping/Traceroute	OS9000E	mpls
- MPLS Traps	OS9000E	mpls
NTP Server	all	base
Server Load Balancing – Weight Round Robin	OS6850/OS9000/OS9000E	base
Hashing Control	OS6850/OS6855/OS9000/OS9000E	base
Source Learning		
- Disable MAC learning per VLAN	OS6400/OS6855-U24X/OS9000E	base
- Disable MAC learning per port	all	base
VRF		
- BFD Support	OS9000E/OS6855-U24X	base
- VRRP Support	OS9000E/OS6855-U24X	base
- Switch Authentication (ASA)	OS9000E/OS6855-U24X	base
- Switch Access and Utilities	OS9000E/OS6855-U24X	base
- Qos Enhancements	OS9000E/OS6855-U24X	base
- UDP/DHCP Relay	OS9000E/OS6855-U24X	base
Ported features for OS9000E		

Feature	Platform	Software Package
- BFD	OS9000E	base
- Configure more than one sFlow receiver	OS9000E	base
- G.8032 Ethernet Ring Protection	OS9000E	base
- IPsec Support for IPv6	OS9000E	base/encrypt
- IPsec Support for OSPF3	OS9000E	base/encrypt
- IPsec Support for RIPng	OS9000E	base/encrypt
- IPv6 Unique Local IPv6 Unicast	OS9000E	base
- IPv6 Scoped Multicast Addresses	OS9000E	base
- Pause Control	OS9000E	base

Previous Release Software Supported

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

Feature/Enhancement Summary

Feature	Platform	Software Package
10Km Stacking	OS6855-U24X	base
31-bit Network Mask Support	all	base
802.1AB MED Extensions	all	base
802.1Q	all	base
802.1Q 2005 (MSTP)	all	base
Access Guardian		base
- 802.1x Device Classification	all	base
- 802.1x RADIUS Failover	all	base
- Captive Portal	all	base
- Captive Portal Web Pages	all	base
- Host Integrity Check (HIC)	6400/6850/6855	base
- User Network Profiles (UNP)	all	base
- QoS Policy Lists	6400/6850/6855	base
Access Control Lists (ACLs)	all	base
- ACLs for IPv4	all	base
- ACLs for IPv6	all	base
- ACL & Layer 3 Security	all	base
- ACL Manager (ACLMAN)	all	base
Account & Password Policies	all	base
ARP Defense Optimization	all	base
ARP Poisoning Detect	all	base
Authenticated Switch Access	all	base
Authenticated VLANs	OS6400/OS6850/OS6855/OS9000	base
Automatic VLAN Containment (AVC)	all	base
Auto-Qos Prioritization of IP Phone Traffic	all	base
Auto-Qos Prioritization of NMS Traffic	all	base
Bi-Directional Forwarding Detection (BFD)	OS6850/OS6855/OS9000/OS9000E	base
BGP Graceful Restart	OS6850/OS6855/OS9000/9000E	advanced routing
BGP4	OS6850/OS6855/OS9000/9000E	advanced routing
BPDU Shutdown Ports	all	base
Command Line Interface (CLI)	all	base
DDM	all	
DHCP		
- Option-82	all	base
- Option 82 – Port-based format	OS6400/OS6850/OS6855	base
- DHCP Relay	all	base
- DHCP Snooping	all	base
- DHCP Snooping Option-82 Data	all	base

Feature	Platform	Software Package
Insertion Format		
DNS Client	all	base
DSCP Range Condition	all	base
DVMRP	OS6850/OS6855/OS9000/OS9000E	advanced routing
Dynamic VLAN Assignment (Mobility)	all	base
Ethernet Ring Protection (G.8032)	**all**	**base**
- Ethernet Ring Protection (ERP) - Shared VLAN	all	base
Ethernet Services		
- L2 Tunneling Enhancements	all	base
- Egress Rate Limiting	OS6400/OS6855-U24X/OS9000E	base
ECMP RIP Support	**OS6850/OS6855/OS9000/9000E**	**base**
- Support for up to 16 paths	OS6850/OS9000/OS9000E	base
End User Partitioning	all	base
Ethernet Interfaces	all	base
Ethernet OAM	**all**	**base**
- Ethernet OAM 802.3ah – EFM	all	base
Flood/Storm Control	all	base
Generic Routing Encapsulation (GRE)	all	base
GVRP	all	base
Hashing Control	OS6850/OS6855/OS9000/OS9000E	base
Health Statistics	all	base
HTTP/HTTPS Port Configuration	all	base
IGMP Multicast Group Configuration Limit	OS6400/OS6850/OS6855/OS9000	base
IGMP Relay - Forward to Specific Host in L3 Environment	OS6850/OS9000/OS9000E	base
Interface Admin Down Warning	OS6400/OS6850/OS6855	base
Interswitch Protocols (AMAP)	All	base
IPMVLAN Multicast Group Overlapping	**all**	**base**
- Group Address and Mask	OS6400/OS6850/OS6855	base
IPMS Flood Unknown Option	all	base
IPsec Support for IPv6	OS6850//OS6855/OS9000/OS9000E	base / encrypt
IPsec Support for OSPF3	OS6850/OS6855/OS9000/OS9000E	base / encrypt
IPsec Support for RIPng	OS6850/OS6855/OS9000/OS9000E	base / encrypt
IPv6		
-Unique Local IPv6 Unicast Addresses	OS6850/OS6855/OS9000/OS9000E	advanced routing
-IPv6 Scoped Multicast Addresses	OS6850/OS6855/OS9000/OS9000E	advanced routing
-IPv6 Multicast Routing	OS6850/OS6855/OS9000/OS9000E	advanced routing
-IPv6 Multicast Switching (MLD)	all	base

Feature	Platform	Software Package
-IPv6 Multicast Switching (Proxying)	all	base
- IPv6 Client and/or Server Support	all	base
- IPv6 Routing	OS6850/OS6855/OS9000/OS9000E	base
IP DoS Filtering	all	base
IP MC VLAN – Support for multiple sender ports	all	base
IP Multinetting	all	base
IP Route Map Redistribution	all	base
IP-IP Tunneling	all	base
IPv4 Multicast Switching (IPMS)	all	base
IPv4 Multicast Switching (Proxying)	all	base
IPv4 Routing	all	base
IS-IS	OS6850/OS9000/OS9000E	advanced routing
ISSU	OS9000E	base
L2 Static Multicast Address	all	base
L4 ACLs over IPv6	all	base
Learned MAC Address Notificaton	all	base
Learned Port Security (LPS)	all	base
Link Aggregation (static & 802.3ad)	all	base
MAC Address Mode	OS9000/OS9000E	base
Mac Authentication for Supplicant/Non-Supplicant	all	base
MAC Retention	OS6400/OS6850/OS6855-U24X	base
Multiple Virtual Routing & Forwarding (Multiple VRF)	OS9000E/OS6855U24X	base
MPLS		
- VPLS Support	OS9000E	mpls
- MPLS Static Fast Re-Route	OS9000E	mpls
- MPLS License	OS9000E	mpls
- MPLS OAM-LSP Ping/Traceroute	OS9000E	mpls
- MPLS Traps	OS9000E	mpls
Network Time Protocol (NTP)		
- Client	all	base
- Server	all	base
OSPFv2	OS6850/OS6855/OS9000/9000E	advanced routing
OSPFv3	OS6850/OS6855/OS9000/9000E	advanced routing
Pause Control/Flow Control	all	base
Port Mapping – Unknown Unicast Flooding	all	base
Partitioned Switch Management	all	base
Pause Control/Flow Control	all	base
Per-VLAN DHCP Relay	all	base
PIM	OS6850/OS6855/OS9000/9000E	advanced routing

Feature	**Platform**	**Software Package**
PIM-SSM (Source-Specific Multicast)		
Policy Based Mirroring	all	base
Policy Based Routing (Permanent Mode)	all	base
Policy Server Management	all	base
Port Mapping	all	base
Port Mirroring (128:1)	all	base
Port Monitoring	all	base
Port-based Ingress Limiting	all	base
Power over Ethernet (PoE)	OS6400/OS6850/OS6855/OS9000	base
PVST+	all	base
Quality of Service (QoS)	all	base
Quarantine Manager and Remediation	all	base
Redirection Policies (Port and Link Aggregate)	all	base
Remote Port Mirroring	all	base
RIPng	OS6850/OS6855/OS9000/OS9000E	base
RIPv1/RIPv2	all	base
RMON	all	base
Router Discovery Protocol (RDP)	all	base
Routing Protocol Preference	all	base
RRSTP	all	base
Secure Copy (SCP)	all	base
Secure Shell (SSH)	all	base
Server Load Balancing	**OS6400/OS6850/OS9000**	**base**
- WRR	OS6850/OS9000/OS9000E	base
sFlow	all	base
Smart Continuous Switching Hot Swap Management Module Failover Power Monitoring Redundancy	all	base
SNMP	all	base
Software Rollback	all	base
Source Learning	all	base
- Disable MAC learning per VLAN	OS6400/OS6855-U24X/OS9000E	base
- Disable MAC learning per port	all	base
Spanning Tree	all	base
SSH Public Key Authentication	all	base
Switch Logging	all	base
Syslog to Multiple Hosts	all	base
Text File Configuration	all	base
TFTP Client for IPv4	all	base
Traffic Anomaly Detection (Network	OS6850/OS6855/OS9000/OS9000E	base

Feature	Platform	Software Package
Security)		
UDLD	all	base
User Definable Loopback Interface	all	base
User Network Profile (UNP)	all	base
VLAN Stacking and Translation	all	base
VLAN Stacking Eservices	all	base
VLANs	all	base
VRF – Multiple VRF Routing and Forwarding	OS9000E/OS6850-U24X	base
- BFD Support	OS9000E/OS6855-U24X	base
- VRRP Support	OS9000E/OS6855-U24X	base
- Switch Authentication (ASA)	OS9000E/OS6855-U24X	base
- Switch Access and Utilities	OS9000E/OS6855-U24X	base
- Qos Enhancements	OS9000E/OS6855-U24X	base
- UDP/DHCP Relay		
VRRP Global Commands	OS6850/OS6855/OS9000/OS9000E	base
VRRPv2	OS6850/OS6855/OS9000/OS9000E	base
VRRPv3	OS6850/OS6855/OS9000/OS9000E	base
Web-Based Management (WebView)	all	base
Webview/SNMP support for BGP IPv6 Extensions	OS6850/OS6855/OS9000/OS9000E	advanced routing
Windows Vista for WebView	all	base

Feature Descriptions

10Km Stacking

The OS6855-U24X supports stacking a maximum of four chassis into a virtual chassis using SFP+ fiber transceivers or directly attached copper SFP+ cables . A distance of up to 10Km is supported using the iSFP-10G-LR fiber transceiver.

802.1AB MED Extensions

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

31-Bit Network Mask Support

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

802.1Q

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.

Access Guardian

802.1x Radius-down Fail-Open

Allows users to be moved to a specified profile when the RADIUS server is not available. This feature is supported for 802.1x and MAC-based authentication, but not for users being authenticated by captive-portal. Users classified through the auth-server-down policy are flagged for re-authentication when the authentication server becomes reachable.

Captive Portal

Captive Portal authentication is a configurable option within Access Guardian that allows Web browser clients to authenticate through the switch using 802.1x or MAC authentication via a RADIUS server. When the Captive Portal option is invoked, a Web page is presented to the user device to prompt the user to enter login credentials. If authentication returns a VLAN ID, the device is assigned to that VLAN. If a VLAN ID is not returned or authentication fails, a separate Captive Portal policy then determines the network access control for the supplicant or non-supplicant.

Captive Portal Web Pages

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

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

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

Captive Portal Browser Support

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

Platforms Supported	Web Browser Supported	Java Version
Windows XP	IE6, IE7, FireFox2 and FireFox3	Java 1.6 update 5 through 12
Windows Vista	IE7, Firefox2 and Firefox3	Java 1.6 update 5 through 12
Linux	Firefox2 and Firefox3	Java 1.6 update 5 through 12

Host Integrity Check (HIC)

Host Integrity Check (HIC) is a mechanism for verifying the compliance of an end user device when it connects to the switch. Configurable HIC policies are used to specify, evaluate, and enforce network access requirements for the host. For example, is the host running a required version of a specific operating system or anti-virus software up to date.

The Access Guardian implementation of HIC is an integrated solution consisting of switch-based functionality, the InfoExpress compliance agent (desktop or Web-based) for the host device, and interaction with the InfoExpress CyberGatekeeper server and Policy Manager. The switch-based functionality is provided through the configuration of a User Network Profile (UNP), which contains a configurable HIC attribute.

NOTE: Minmum ASIC versions are required for HIC support as noted below. Use the '**show ni'** command documented in the **Supported Hardware/Software Combinations** section to verify the ASIC version.

Platform	ASIC Version Required
6850/6855	B2
6400/6855-U24X	A0

Host Integrity Check Platform and Browser Support

The HIC switch-based functionality interacts with compliance agents and the CyberGatekeeper server from InfoExpress. The compliance products consist of a desktop and Web-based agent. The following table provides platform and browser support information for both types of agents:

Compliance Agent	Platforms Supported	Web Browser Supported
Desktop	Windows Vista, XP, 2003, 2000 Linux (Red Hat and SUSE Dists.)	N/A

Compliance Agent	Platforms Supported	Web Browser Supported
Web-based	Windows Vista, XP, 2003, 2000	IE versions 6 and 7 Firefox 2.x, Firefox 3.x Java 1.6 update 5 through 12

Refer to the InfoExpress documentation for information about how to configure the CyberGatekeeper server and other related products.

User Network Profile (UNP)

A User Network Profile (UNP) defines network access controls for one or more user devices. Each device that is assigned to a specific profile is granted network access based on the profile criteria, instead of on an individual MAC address, IP address, or port. Assigning users to a profile provides greater flexibility and scalability across the network. Administrators can use profiles to group users according to function. All users assigned to the same UNP become members of that profile group. The UNP then determines what network access resources are available to a group of users, regardless of source subnet, VLAN or other characteristics.

A UNP is a configurable option of Access Guardian device classification policies and consists of the following attributes:

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

A UNP is a configurable option of Access Guardian device classification policies. A policy may also include 802.1X, MAC, or Captive Portal (Web-based) authentication to provide more granular control of the profile.

One of the attributes of a User Network Profile (UNP) specifies the name of a list of QoS policy rules. This list is applied to a user device when the device is assigned to the user profile. Using policy lists allows the administrator to associate a group of users to a set of QoS policy rules.

A default policy list exists in the switch configuration. Rules are automatically added to this list when the rule is created. A rule can belong to multiple policy lists. As a result, the rule remains a member a of the default list even when it is subsequently assigned to additional lists. The user does have the option to exclude the rule from the default list to preserve system resources.

Up to 13 policy lists (including the default list) are supported per switch. Only one policy list per UNP is allowed, but a policy list can be associated with multiple profiles.

Access Control Lists (ACLs)

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

In general, the types of ACLs include:

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

Access Control Lists (ACLs) for IPv6

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

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

Note the following when using IPv6 ACLs:

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

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

ACL Manager

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

This implementation of ACLMAN also provides the following features:

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

Account & Password Policies

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

ARP Defense Optimization

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

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

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

The following table provides the platforms and browser support information for AVLAN web authentication:

Platforms Supported	Web Browser Supported	Java Version
Windows 2000	IE6	Java 1.6 update 5 through 12
Windows XP	IE6, IE7, FireFox2, FireFox3, Netscape 9.0	Java 1.6 update 5 through 12
Windows Vista	IE7, Firefox3, Netwscape 9.0	Java 1.6 update 5 through 12
Linux	Netscape 4.75 and later	--
MAC OS 10.5	Safari 3.0.4	Java 12.0

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.

Bi-Directional Forwarding Detection (BFD)

Bidirectional Forwarding Detection (BFD) is a hello protocol that can be configured to interact with routing protocols for the detection of path failures and can reduce the convergence time in a network. BFD is supported with the following Layer 3 protocols: BGP, OSPF, VRRP Tracking and Static Routes.

When BFD is configured and enabled, BFD sessions are created and timers are negotiated between BFD neighbors. If a system does not receive a BFD control packet within the negotiated time interval, the neighbor system is considered down. Rapid failure detection notices are then sent to the routing protocol, which initiates a routing protocol recalculation. This process can reduce the time of convergence in a network.

BGP4

The Border Gateway Protocol (BGP) is an exterior routing protocol that guarantees the loop-free exchange of routing information between autonomous systems. The Alcatel-Lucent implementation supports BGP version 4 as defined in RFCs 1771/4271, 2439, 3392, 2385, 1997, 4456, 3065, 4273 and 4486.

The Alcatel-Lucent implementation of BGP is designed for enterprise networks, specifically for border routers handling a public network connection, such as the organization’s Internet Service Provider (ISP) link. Up to 65,000 route table entries and next hop routes can be supported by BGP.

BGP IPv6 Extensions

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

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

The feature includes Webview and SNMP support.

BGP Graceful Restart

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

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.

DDM - Digital Diagnostic Monitoring

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

- Temperature
- Supply Voltage
- Current
- Output Power
- Input Power

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

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

Detect ARP Poisoning

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

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

DHCP Relay

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

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

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

DHCP Relay Agent Information Option

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

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

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

- The OmniSwitch enhances the Option 82 capability by allowing the 'interface alias' to be inserted into the Circuit ID and Remote ID suboptions of the Option-82 field

DHCP Snooping

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

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

Additional DHCP Snooping functionality includes the following:

- **Layer 2 DHCP Snooping—**Applies DHCP Snooping functionality to bridged DHCP client/server broadcasts without using the relay agent or requiring an IP interface on the client/server VLAN.
- **IP Source Filtering—**Restricts DHCP Snooping port traffic to only packets that contain the client source MAC address and IP address obtained from the DHCP lease information. The DHCP Snooping binding table is used to verify the client lease information for the port that is enabled for IP source filtering.
- **Rate Limiting—**Limits the number of DHCP packets on a port. This functionality is provided using the QoS application to configure ACLs for the port.
- **User-Configurable Option 82 Suboption Format—**Allows the user to specify the type of information (switch base MAC address, system name, or user-defined string) that is inserted into the Circuit ID and Remote ID suboptions of the Option-82 field. This functionality only applies when DHCP Snooping Option-82 Data Insertion is enabled.

DNS Client

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

Dynamic VLAN Assignment (Mobility)

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

DVMRP

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

End User Partitioning (EUPM)

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

Ethernet Interfaces

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

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

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

Ethernet OAM

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

Ethernet First Mile (EFM)

IEEE 802.3ah, defining Ethernet in the access networks that connects subscribers to their immediate service provider. EFM, EFM-OAM and LINKOAM refers to IEEE 802.3ah standard.

LINK OAM (operation, administration, and maintenance) is a tool which monitors Layer-2 link status on the network by sending OAM protocol data units (OAMPDUs) between the network devices. OAMPDUs contain control and status information used to monitor, test and troubleshoot OAM-enabled links. By enabling LINK OAM on switch ports, network administators can monitor the link-related issues on the first mile. LINK OAM provides network administrators the ability to monitor link performance, remote fault detection and remote loopback control.

Ethernet Ring Protection (ERP) – G.8032

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

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

ERP – Overlapping Protected VLANs on a Single Node

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

Ethernet Services

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

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

- Ethernet service-based approach that is similar to configuring a virtual private LAN service (VPLS).
- Ingress bandwidth sharing across User Network Interface (UNI) ports.
- Ingress bandwidth rate limiting on a per UNI port, per CVLAN, or CVLAN per UNI port basis.
- CVLAN (inner) tag 802.1p-bit mapping to SVLAN (outer) tag 802.1p bit.
- CVLAN (inner) tag DSCP mapping to SVLAN (outer) tag 802.1p bit.
- Profiles for saving and applying traffic engineering parameter values.

Ethernet Services - Egress rate limiting

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

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

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

Ethernet Services - Tunneling L2 Protocols

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

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

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. Up to 32 UDP instances can be configured.

Generic Routing Encapsulation

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

GVRP

The GARP VLAN Registration Protocol (GVRP), a protocol compliant with 802.1Q, dynamically learns and further propagates VLAN membership information across a bridged network. GVRP dynamically maintains and updates the registration and de-registration of VLANs and prunes unnecessary broadcast and unicast traffic. Through propagation of GVRP information, a device is continuously able to update its knowledge of the set of VLANs that currently have active members and of the ports through which those members can be reached. With GVRP, a single switch is manually configured with all the desired VLANs for the network, and all other switches on the network dynamically learn those VLANs. An end station can be plugged into any switch and can be connected to its desired VLAN. However, for end stations to make use of GVRP, they need Network Interface Cards (NIC) aware of GVRP.

Hashing Control

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

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

Default Hashing Mode and Recommendations

Platform	Default Hashing Mode
9000/9000E	Extended
6400/6850/6855	Brief

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

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

The hash control mode also impacts the fabric load balancing for chassis-based products. It is not recommended to set brief hashing mode on chassis-based products

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.

IGMP Multicast Group Configuration Limit

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

IGMP Relay - Relay IGMP Packets to Specific Host

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

Interface Admin Down Warning

The user can enable/disable the display of a confirmation prompt before an interface is administratively disabled to prevent a user from inadvertantly issuing an “admin down” command for an interface(s). This feature is disabled by default.

IP Multicast Flood Unknown

When this feature is enabled, multicast packets are flooded on the VLAN until the multicast group membership table is updated, they are then forwarded based on the multicast group membership table.

IPMVLAN Multicast Group Overlapping

Different ISPs may use the same multicast group addresses. To remedy this, a user can configure the same multicast address on different IP Multicast VLANs (IPMVLAN). A common use case will be a network where each receiver port is only configured for one IPMVLAN. A user can define the mapping between an IPMVLAN and a customer VLAN ID (c-tag) to be used in the c-tag translation rule. Additionally, this feature allows a mask to be specified.

IPsec Support for IPv6

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

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

Note: This is a licensed feature and requires that a license file be installed on the switch. Refer to the current price list for ordering information.

IPv6 - Globally Unique Local Unicast Addresses

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

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

IPv6 – Scoped Multicast Addresses

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

IP/IP Tunneling

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

IP Multicast VLAN

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

Inlcudes support for multiple sender ports.

Interswitch Protocol (AMAP)

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

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

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

IPv4 Support

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

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Bootstrap Protocol (BOOTP)/Dynamic Host Configuration Protocol (DHCP)
- Simple Network Management Protocol (SNMP)
- Telnet - Client and server
- File Transfer Protocol (FTP) – Client and server
- Address Resolution Protocol (ARP)

- Internet Control Message Protocol (ICMP)
- RIP I / RIP II
- Static Routes

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

IPv6 Support

IPv6 (documented in RFC 2460) is designed as a successor to IPv4 and is supported on the OmniSwitch 6850, 6855 and 9000/9000E. The changes from IPv4 to IPv6 fall primarily into the following categories:

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

IP DoS Filtering

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

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

IP Multicast Switching (IPMS)

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

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

Both IGMP version 3 (IGMPv3), which handles forwarding by source IP address and IP multicast destination, and IGMP version 2 (IGMPv2), which handles forwarding by IP multicast destination address only, are supported.

IP Multicast Switching (IPMS) - Proxying

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

IP Multinetting

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

IP Route Map Redistribution

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

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

IS-IS

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

In-Service Software Upgrade (ISSU)

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

- Switches running an '**R**##' build, such as 6.4.2.123.R01 **do not** support ISSU upgrades. The switch must first be upgraded to an '**S##**' build such as 6.4.2 .123.**S01**.

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

- ISSU patches are only supported within the same 'S##' branch. For example, if a switch is running 6.4.2.123.S01 then only 6.4.2.###.S01 images can used to perform an ISSU patch. If a switch is running 6.4.2.234.S02 then only 6.4.2.###.S02 images can used to perform an ISSU patch.

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

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

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

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

L2 Static Multicast Addresses

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

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

Learned Port Security (LPS)

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

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

LPS has the following limitations:

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

Learned MAC Address Notification

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

Link Aggregation (static & 802.3ad)

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

- **Scalability (OS6400/6850/6855).** You can configure up to 32 link aggregation groups that can consist of 2, 4, or 8 Ethernetports.
- **Scalability (OS9000/OS9000E).** You can configure up to 128 link aggregation groups that can consist of 2, 4, or 8 Ethernetports.
- **Reliability.** If one of the physical links in a link aggregate group goes down, the link aggregate group can still operate.
- **Ease of Migration.** Link aggregation can ease the transition from a Gigabit Ethernet backbone to a 10 Gigabit Ethernet backbone.
- **Interoperability with Legacy Switches.** Static link aggregation can interoperate with OmniChannel on legacy switches.

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

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

Number of ports in group	Maximum number of groups
2	128
4	64
8	32

Multiprotocol Label Switching (MPLS)

Multiprotocol Label Switching (MPLS) is a label switching technology that provides the ability to set up connection-oriented paths over a connectionless IP network. MPLS sets up a specific path for a sequence of packets. The packets are identified by a label inserted into each packet.

This implementation of MPLS provides the network architecture that is needed to set up a Virtual Private LAN Service (VPLS). VPLS allows multiple customer sites to transparently connect through a single bridging domain over an IP/MPLS-based network.

The MPLS architecture provided is based on the Label Distribution Protocol (LDP). The LDP consists of a set of procedures used by participating Label Switching Routers (LSRs) to define Label Switched Paths (LSPs), also referred to as MPLS tunnels. These tunnels provide the foundation necessary to provision VPLS.

MPLS Software Licensing Requirement. The MPLS feature, including the VPLS application, requires the purchase of an Alcatel-Lucent software license. The licenses are available through the Alcatel-Lucent Software License portal.

VPLS Support

A Virtual Private LAN Service (VPLS) is a Virtual Private Network (VPN) technology that allows any-to-any (multipoint) connectivity. The provider network emulates a LAN by connecting all the remote customer sites at the edge of the provider network to a single bridged LAN. A full mesh of pseudo-wires (PW) is established to form a VPLS.

A VPLS-capable network consists of Customer Edges (CE), Provider Edges (PE), and a core MPLS network. The IP/MPLS core network interconnects the PEs but does not participate in the VPN functionality. Traffic is simply switched based on the MPLS labels.

This implementation of VPLS makes use of a service-based architecture that provides the following logical entities that are required to provision a service:

- **Customers (subscribers).** An account is created for each customer and assigned an ID. The customer ID is required and associated with the service at the time the service is created.
- **Service Access Points (SAPs).** Each subscriber service type is configured with at least one SAP. A SAP identifies the point at which customer traffic enters the service.
- **Service Distribution Points (SDPs)**. A SDP provides a logical point at which customer traffic is directed from one PE to another PE through a one-way service tunnel.

MPLS Static Fast Re-Route

MPLS forwarding is performed by routers called Label Switching Routers (LSRs). A Label Switched Path (LSP) is a path through one or more LSRs.

There are two types of LSPs that are configurable using MPLS:

- **Static LSPs**. A Static LSP specifies a statically defined path of LSRs. Configuration of label mappings and MPLS actions is required on each router that will participate in the static path. No signaling protocol, such as the Label Distribution Protocol (LDP), is required, and there is no dependence on a gateway protocol topology or local forwarding table. Static LSPs are able to cross an Autonomous System (AS) boundary.
- **Signaled LSP**. The LSPs are set up using a signaling protocol, such as LDP. The signaling protocol allows the automatic assignment of labels from an ingress router to the egress router. Signaling is triggered by the ingress router, therefore configuration is only required on this router. A signaled LSP is confined to one gateway protocol area and, therefore, cannot cross an AS boundary.

In addition to static LSPs, a static Fast Reroute (FRR) feature is available that allows the configuration of backup static LSP tunnels. FRR uses these backup tunnels to provide alternate routes in the event an LSP goes down.

MPLS OAM-LSP Ping/Traceroute

When an MPLS Label Switched Path (LSP) fails to deliver customer traffic, the failure is not always detected by the MPLS control plane. To assist users with detecting and isolating traffic problems, such as “black holes” or incorrect routing, the following MPLS OAM (Operations, Administration, and Maintenance) tools are available:

- LSP Ping to perform connectivity checks.
- LSP Traceroute to perform hop-by-hop fault localization and path tracing.

LSP Ping and Traceroute are used to verify that packets associated with a particular Forwarding Equivalence Class (FEC) actually end their MPLS path on a Label Switching Router (LSR) that is an Egress LSR for that FEC.

MPLS Traps

The OmniSwitch AOS implementation of MPLS generates the following SNMP traps.

mplsXCup
mplsXCdown
vRtrMplsStateChange
vRtrMplsIfStateChange
vRtrMplsLspUp
vRtrMplsLspDown
vRtrLdpInstanceStateChange
vRtrLdpGroupIdMismatch

svcStatusChanged
sapStatusChanged
sdpBindStatusChanged
sdpStatusChanged
sapPortStateChangeProcessed
sdpBindStateChangeProcessed
sdpKeepAliveProbeFailure
sdpKeepAliveStarted
sdpKeepAliveStopped

Multiple Virtual Routing and Forwarding (Multiple-VRF)

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

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

- Multiple routing instances within the same physical switch. Each VRF instance is associated with a set of IP interfaces and creates and maintains independent routing tables. Traffic between IP interfaces is only routed and forwarded within those interfaces/routes that belong to the same VRF instance.
- Multiple instances of IP routing protocols, such as static, RIP, IPv4, BGPv4, and OSPFv2 on the same physical switch. An instance of each type of protocol operates within its own VRF instance.

- The ability to use duplicate IP addresses across VRF instances. Each VRF instance maintains its own IP address space to avoid any conflict with the service provider network or other customer networks.
- Separate IP routing domains for customer networks. VRF instances configured on the Provider Edge (PE) are used to isolate and carry customer traffic through the shared provider network.

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

VRF - Qos Enhancements

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

VRF - Switch Authentication Enhancement

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

VRF - Switch Access and Utilities

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

VRF - VRRP

Enhances VRRP making it VRF aware. Allows for the configuration of independent VRRP instances in multiple VRFs.

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

VRF – UDP/DHCP Relay

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

- VRF-Aware commands. These commands are allowed in any of the VRF instances configured in the switch. The settings in one VRF are independent of the settings in another VRF. Command parameters are visible and configurable within the context of any VRF.

- Global commands. These commands are supported only in the default VRF, but are visible and applied to all VRF instances configured in the switch. This command behavior is similar to how command parameters are applied in the per-VLAN DHCP Relay mode. For example, the maximum hops value configured in the default VRF is applied to all DHCP Relay agents across all VRF instances. This value is not configurable in any other VRF instance.

- Default VRF commands. These commands are supported only in the default VRF and are not applied to any other VRF instance configured in the switch. For example, per-VLAN mode, DHCP Snooping, and boot-up commands fall into this category.

Refer to the “Configuring Multiple VRF” chapter in the OmniSwitch AOS Release 6 Configuration Guide for a list of UDP/DHCP Relay VRF related commands.

Note: Refer to the “Configuring Multiple VRF” chapter in the OmniSwitch AOS Release 6 Configuration Guide for a list of VRF supported features and commands.

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

Pause Control/Flow Control

PAUSE frames are used to pause the flow of traffic between two connected devices when traffic congestion occurs. PAUSE frame flow control provides the ability to configure whether or not the switch will transmit and/or honor PAUSE frames on an active interface. This feature is only supported on interfaces configured to run in full-duplex mode.

In addition to configured PAUSE frame flow control settings, this feature also works in conjunction with auto-negotiation to determine operational transmit/receive settings for PAUSE frames between two switches. Note that the configured PAUSE frame flow control settings are overridden by the values that are determined through auto-negotiation.

End-to-end flow control is supported on OmniSwitch 6400, 6850, and 6855 switches running in standalone mode. When working in stack mode, these switches will honor received pause messages on any port of any stack. In the case of an OmniSwitch chassis, received pause frames will be honored and processed.

To enable end to end flow control on 48-port standalone OmniSwitch 6400 and 6850 switches, a dedicated VLAN must be configured *and* RX/TX pause enabled. In the case of 24-port standalone switches, enabling RX/TX pause is sufficient.

Port Mapping – Unknown Unicast Flooding

By default, unknown unicast traffic is flooded to the user ports of a port mapping session from all the switch ports, not just the network ports for the session. There is now a port mapping option to enable or disable unknown unicast flooding from network ports to user ports.

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).

NTP Server

Enhances the NTP functionality to allow the OmniSwitch to act as an NTP server. The OmniSwitch software by default will be able to respond to NTP client requests, and establish a client/server peering relationship. With the server cli commands now enabled, the Omniswitch can now also establish an active peering relationship with another server, enable broadcast server functionality, disable a given IP for NTP and employ MD5 authentication for clients and active peers.

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.

PIM-SM/PIM-DM/PIM-SSM

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

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

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

Policy Server Management

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

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

Policy Based Routing (Permanent Mode)

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

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

Port Mapping (Private VLANs)

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

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

Port Monitoring

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

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

Power over Ethernet (PoE)

The Power over Ethernet (PoE) software is supported on the various OmniSwitch platforms. 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.

PVST+ Interoperability

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

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

Quality of Service (QoS)

Alcatel-Lucent’s QoS software provides a way to manipulate flows coming through the switch based on user-configured policies. The flow manipulation (generally referred to as Quality of Service or QoS) may be as simple as allowing/denying traffic, or as complicated as remapping 802.1p bits from a Layer 2 network to ToS values in a Layer 3 network. QoS can support up to 2048 policies and it is hardware-based on the first packet. OmniSwitch 6850/9000/9000E switches 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
- **Addded support for DSCP Ranges**
- **Policy Based Routing (PBR)—**includes policies for redirecting routed traffic
- **Access Control Lists (ACLs)—**ACLs are a specific type of QoS policy used for Layer 2, Layer 3/4, and multicast filtering.

Auto-Qos Prioritization for NMS Traffic

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

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

Auto-Qos Prioritization on IP Phones

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

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

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

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

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

BPDU Shutdown Ports

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

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

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

Policy-Based Mirroring

This feature enhances the current port mirroring functionality on the OmniSwitch. It allows policies to be configured to determine when traffic should be mirrored based on policies rather than being restricted to a specified port. The following policies can be configured:

- Traffic between 2 ports

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

Policy-Based Mirroring limitations:

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

Ingress and Egress Bandwidth Shaping

Bandwidth shaping is configured on a per port basis by specifying a maximum bandwidth value for ingress and egress ports. However, on the OmniSwitch 6850 and 9000/9000E switches, configuring minimum and maximum egress bandwidth is supported on a per COS queue basis for each port.

Quarantine Manager and Remediation (QMR)

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

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

- **Remediation server and exception subnet group.** This is a reserved QoS network group, called “alaExceptionSubnet”, that is configured with the IP address of a remediation server and any subnets to which a quarantined client is allowed access. The quarantined client is redirected to the remediation server to obtain updates and correct its quarantined state.
- **Remediation server URL.** This is the URL for the remediation server. Note that this done in addition to specifying the server IP address in the “alaExceptionSubnet” network group.
- **Quarantined Page.** When a client is quarantined and a remediation server URL is not configured, QMR can send a Quarantine Page to notify the client of its quarantined state.
- **HTTP proxy port group.** This is a known QoS service group, called “alaHTTPProxy”, that specifies the HTTP port to which quarantined client traffic is redirected for remediation. The default HTTP port used is TCP 80 and TCP 8080.

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

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

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

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

Remote Port Mirroring (802.1Q Based)

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

- There must not be any physical loop present in the remote port mirroring VLAN.
- Spanning Tree must be disabled for the remote port mirroring VLAN.
- BPDU mirroring will be disabled by default on OS6400/6850/68555 switches.
- BPDU mirroring will be disabled by default on all OS9000s with B2 revision ASICs. (Contact Service and Support to enable)
- BPDU mirroring will be enabled by default on all OS9000s with A0/A1 revision ASICs.
- Source learning must be disabled or overridden on the ports belonging to the remote port mirroring VLAN on the intermediate and destination switches.
- The QoS redirect feature can be used to override source learning.

RIPv1/RIPv2

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

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

RIPng

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

RIP Timer Configuration

- Update —The time interval between advertisement intervals.
- Invalid—The amount of time before an active route expires and transitions to the garbage state.

- Garbage—The amount of time an expired route remains in the garbage state before it is removed from the RIB.
- Holddown—The amount of time during which a route remains in the hold-down state.

Redirect Policies (Port and Link Aggregate)

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

RMON

Remote Network Monitoring (RMON) is an SNMP protocol used to manage networks remotely. RMON probes can be used to collect, interpret, and forward statistical data about network traffic from designated active ports in a LAN segment to an NMS (Network Management System) application for monitoring and analyzing without negatively impacting network performance. RMON software is fully integrated in the software to acquire statistical information.

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

Router Discovery Protocol (RDP)

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

Routing Protocol Preference

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

RRSTP

Ring Rapid Spanning Tree Protocol (RRSTP) is complimentary to either the Rapid Spanning Tree (RSTP) or the Multiple Spanning Tree Protocol (MSTP) but is designed to enhance convergence time in a ring configuration when a link failure occurs. Note that RRSTP is supported only in a ring topology where switches are connected point to point. In addition, there can be no alternate connections for the same instance between any two switches within a ring topology.

RRSTP reduces convergence time by finding the bridge that hosts the alternate (ALT) port and immediately changing the ALT port state to forwarding without altering the port state. This process quickly enables the data path. The RRSTP frame travels from the point of failure to the ALT port in both directions. The MAC addresses corresponding to the ports in the ring are flushed to make the data path convergence time much faster. While RRSTP is already reacting to the loss of connectivity, the standard BPDU carrying the information about the link failure is processed in normal fashion at each hop. When this BPDU reaches the bridge whose ALT port is now in the "ALT FWD" state, due to RRSTP frame processing, it updates the state of the two ports in the ring as per the STP standard.

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

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

Using SLB clusters can provide cost savings (costly hardware upgrades can be delayed or avoided), 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).

Server Load Balancing - WRR

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

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).

Up to two sFlow receivers can be configured.

Smart Continuous Switching – Stackable Products

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

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

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

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

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

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

Smart Continuous Switching – Chassis-based Products

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

The OS9600 operates with one CMM, the other chassis contain two CMM slots.

If there are two CMM modules in a chassis, 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 a chassis, both fabric modules are normally active. Two CMM modules must be installed in the chassis 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 a chassis, then there is a single management processor, but there is no "secondary" CMM. Hardware or software failures in the CMM may result in a system reboot. The system fabric capacity on a system with one CMM is only 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. The OmniSwitch supports SNMPv1, SNMPv2, and SNMPv3.

Source Learning

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

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

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

Disable Learning on a per port basis

Provides the option to disable source learning on a per port basis. This feature is only supported on "hardware learning" ports and is not supported on mobile ports, LPS ports or Access Guardian ports. The feature is also supported for Link Aggregation where all ports in the aggregate are set to disable source learning. Configuration of static mac-addresses on such ports is still allowed.

Disable MAC learning on a per VLAN basis

Provides the option to disable source learning for all the ports of a VLAN. This feature is meant to be used on a ring topology where a VLAN only contains two ports.

It is recommended to have only 2 ports in a VLAN that has source learning disabled.

MAC Address Mode

There are two source learning modes available for the OmniSwitch chassis-based systems: 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 chassis increases the table size to 16K per module and up to 64K per system..

Note: The distributed MAC address mode is only supported chassis-based systems.

Software Rollback

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

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

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

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

Spanning Tree

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

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

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

Syslog to Multiple Hosts

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

Switch Logging

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

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

Text File Configuration

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

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

TFTP Client for IPv4

Trivial File Transfer Protocol (TFTP), a client-server protocol, can be used to transfer files between the TFTP server and client. TFTP client functionality on the OmniSwitch is used to download files from or upload files to the TFTP server within a LAN.

Traffic Anomaly Detection (TAD)

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

Network Security provides the following capabilities:

- Real time network traffic monitoring.
- Dynamic anomaly detection.
- Dynamic anomalous port quarantining.

UDLD - Fiber and Copper

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

User Definable Loopback Interface

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

User Network Profile (UNP)

A User Network Profile (UNP) defines network access controls for one or more user devices. Each device that is assigned to a specific profile is granted network access based on the profile criteria, instead of on an individual MAC address, IP address, or port. Assigning users to a profile provides greater flexibility and scalability across the network. Administrators can use profiles to group users according to function. All users assigned to the same UNP become members of that profile group. The UNP then determines what network access resources are available to a group of users, regardless of source subnet, VLAN or other characteristics.

VLANs

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

The VLAN management software handles the following VLAN configuration tasks:

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

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

VRRPv2/VRRPv3

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

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

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

Authentication is not supported.

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

VRRPv2 is available on all supported OmniSwitch platforms in this release.

Global VRRP Configuration

The following capabilities for VRRP2 were added:

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

Web-Based Management (WebView)

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

- IE6, IE7, Firefox 2, Firefox 3 for Windows NT, 2000, 2003, XP, Windows Vista
- Firefox 2.0 for Solaris SunOS 5.10

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

SNMP Traps

The following table provides a list of AOS Release 6.4.4.R01 SNMP traps managed by the switch.

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

No.	Trap Name	Platforms	Description
			neighbors on the same interface with a lower IP address than itself.
22	dvmrpNeighborLoss	all	A 2-way adjacency relationship with a neighbor has been lost. This trap is generated when the neighbor state changes from “active” to “one-way,” “ignoring” or “down.” The trap is sent only when the switch has no other neighbors on the same interface with a lower IP address than itself.
23	dvmrpNeighborNotPruning	all	A non-pruning neighbor has been detected in an implementation-dependent manner. This trap is generated at most once per generation ID of the neighbor. For example, it should be generated at the time a neighbor is first heard from if the prune bit is not set. It should also be generated if the local system has the ability to tell that a neighbor which sets the prune bit is not pruning any branches over an extended period of time. The trap should be generated if the router has no other neighbors on the same interface with a lower IP address than itself.
24	risingAlarm	all	An Ethernet statistical variable has exceeded its rising threshold. The variable’s rising threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
25	fallingAlarm	all	An Ethernet statistical variable has dipped below its falling threshold. The variable’s falling threshold and whether it will issue an SNMP trap for this condition are configured by an NMS station running RMON.
26	stpNewRoot	all	Sent by a bridge that became the new root of the spanning tree.
27	stpRootPortChange	all	A root port has changed for a spanning tree bridge. The root port is the port that offers the lowest cost path from this bridge to the root bridge.
28	mirrorConfigError	-	Unsupported.
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 pesudoCAM 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.

No.	Trap Name	Platforms	Description
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	6400 6850 6850E 6855	Two or more slots claim to have the same slot number.
38	alaStackMgrNeighborChangeTrap	6400 6850 6850E 6855	Indicates whether or not the stack is in loop.
39	alaStackMgrRoleChangeTrap	6400 6850 6850E 6855	Indicates that a new primary or secondary stack is elected.
40	lpsViolationTrap	all	A Learned Port Security (LPS) violation has occurred.
41	alaDoSTrap	all	Indicates that the sending agent has received a Denial of Service (DoS) attack.
42	gmBindRuleViolation	all	Occurs whenever a binding rule which has been configured gets violated.
43	unused	-	-
44	unused	-	-
45	unused	-	-
46	unused	-	-
47	pethPsePortOnOff	-	Indicates if power inline port is or is not delivering power to the a power inline device.
48	pethPsePortPowerMaintenanceStatus	-	Indicates the status of the power maintenance signature for inline power.
49	pethMainPowerUsageOn	-	Indicates that the power inline usage is above the threshold.
50	pethMainPowerUsageOff	-	Indicates that the power inline usage is below the threshold.
51	ospfNbrStateChange	all	Indicates a state change of the neighbor relationship.
52	ospfVirtNbrStateChange	all	Indicates a state change of the virtual neighbor relationship.
53	httpServerDoSAttackTrap	all	This trap is sent to management station(s) when the HTTP server is under Denial of Service attack. The HTTP and HTTPS connections are sampled at a 15 second interval. This trap is sent every 1 minute while the HTTP server detects it is under attack.
54	alaStackMgrDuplicateRoleTrap	6400 6850 6850E 6855	The element identified by alaStack-MgrSlotNINumber detected the presence of two elements with the same primary or secondary role as specified by alaStackMgrChasRole on the stack.
55	alaStackMgrClearedSlotTrap	6400	The element identified by alaStack-

No.	Trap Name	Platfor ms	Description
		6850 6850E 6855	MgrSlotNINumber will enter the pass through mode because its operational slot was cleared with immediate effect.
56	alaStackMgrOutOfSlotsTrap	6400 6850 6850E 6855	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	6400 6850 6850E 6855	The element identified by alaStack MgrSlotNINumber will enter the pass through mode because there are no tokens available to be assigned to this element.
58	alaStackMgrOutOfPassThruSlotsTrap	6400 6850 6850E 6855	There are no pass through slots avail able to be assigned to an element that is supposed to enter the pass through mode.
59	gmHwVlanRuleTableOverloadAlert	all	An overload trap occurs whenever a new entry to the hardware VLAN rule table gets dropped due to the overload of the table.
60	lnkaggAggUp	all	Indicates the link aggregate is active. This trap is sent when any one port of the link aggregate group goes into the attached state.
61	lnkaggAggDown	all	Indicates the link aggregate is not active. This trap is sent when all ports of the link aggregate group are no longer in the attached state.
62	lnkaggPortJoin	all	This trap is sent when any given port of the link aggregate group goes to the attached state.
63	lnkaggPortLeave	all	This trap is sent when any given port detaches from the link aggregate group.
64	lnkaggPortRemove	all	This trap is sent when any given port of the link aggregate group is removed due to an invalid configura tion.
65	pktDrop	all	The pktDrop trap indicates that the sending agent has dropped certain packets (to blocked IP ports, from spoofed addresses, etc.).
66	monitorFileWritten	-	A File Written Trap is sent when the amount of data requested by the user has been written by the port monitoring instance.
67	alaVrrp3TrapProtoError	all	Indicates that a TTL, checksum, or version error was encountered upon receipt of a VRRP advertisement.
68	alaVrrp3TrapNewMaster	all	The SNMP agent has transferred from the backup state to the master state.
69	gmHwMixModeSubnetRuleTableOverloadAlert	all	A subnet overload trap occurs in mixed mode whenever a new entry to the HW subnet rule table gets dropped due to the overload of the table.
70	pethPwrSupplyConflict	all	Power supply type conflict trap.
71	pethPwrSupplyNotSupported	all	Power supply not supported trap.

No.	Trap Name	Platforms	Description
72	lpsPortUpAfterLearningWindowExpiredTrap	all	When an LPS port joins or is enabled after the Learning Window is expired, the MAC address learning on the port will be disabled, and this trap is generated as a notification.
73	vRtrIsisDatabaseOverload	all	This notification is generated when the system enters or leaves the Overload state.
74	vRtrIsisManualAddressDrops	all	Generated when one of the manual area addresses assigned to this system is ignored when computing routes.
75	vRtrIsisCorruptedLSPDetected	all	This notification is generated when an LSP that was stored in memory has become corrupted.
76	vRtrIsisMaxSeqExceedAttempt	all	Generated when the sequence number on an LSP wraps the 32 bit sequence counter
77	vRtrIsisIDLenMismatch	all	Need Desc. A notification sent when a PDU is received with a different value of the System ID Length.
78	vRtrIsisMaxAreaAddrsMismatch	all	A notification sent when a PDU is received with a different value of the Maximum Area Addresses.
79	vRtrIsisOwnLSPPurge	all	A notification sent when a PDU is received with an OmniSwitch systemID and zero age
80	vRtrIsisSequenceNumberSkip	all	When we recieve an LSP is received without a System ID and different contents.
81	vRtrIsisAutTypeFail	all	A notification sent when a PDU is received with the wrong authentication type field.
82	vRtrIsisAuthFail	all	A notification sent when a PDU is received with an incorrent authentication information field.
83	vRtrIsisVersionSkew	all	A notification sent when a a Hello PDU is received from an IS running a different version of the protocol.
84	vRtrIsisAreaMismatch	all	A notification sent when a Hello PDU is received from an IS which does not share any area address.
85	vRtrIsisRejectedAdjacency	all	A notification sent when a Hello PDU is received from an IS, but does not establish an adjacency due to a lack of resources.
86	vRtrIsisLSPTooLargeToPropagate	all	A notification sent when an attempt to propagate an LSP which is larger than the dataLinkBlockSize for a circuit.
87	vRtrIsisOrigLSPBufSizeMismatch	all	A notification sent when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for the originating L1LSP BufferSize or originating L2LSPBufferSize respectively. Also when a Level 1 LSP or Level2 LSP is received containing the originating LSPBufferSize option and the value in the PDU option field does not match the local value for originating L1LSP BufferSize or originatingL2LSP BufferSize respectively.

No.	Trap Name	Platforms	Description
88	vRtrIsisProtoSuppMismatch	all	A notification sent when a non-pseudonode segment 0 LSP is received that has no matching protocols supported.
89	vRtrIsisAdjacencyChange	all	A notification sent when an adjacency changes state, entering or leaving state up. The first 6 bytes of the vRtrIsisTrapLSPID are the SystemID of the adjacent IS.
90	vRtrIsisCircIdExhausted	all	A notification sent when ISIS cannot be started on a LAN interface because a unique circId could not be assigned due to the exhaustion of the circId space.
91	vRtrIsisAdjRestartStatusChange	all	A notification sent when an adjancency's graceful restart status changes.
92	dot1agCfmFaultAlarm	all	A MEP has lost contact with one or more MEPs. A notification (fault alarm) is sent to the management entity with the OID of the MEP that has detected the fault.
93	Unused	all	-
94	lldpRemTablesChange	all	A lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes.
95	chassisTrapsPossibleDuplicateMac	6400 6850 6850E 6855	The old PRIMARY element cannot be detected in the stack. There is a possiblity of a duplicate MAC address in the network
96	unused	all	-
97	alaPimInvalidRegister	all	An alaPimInvalidRegister notification signifies that an invalid PIM Register message was received by this device
98	alaPimInvalidJoinPrune	all	A alaPimInvalidJoinPrune notification signifies that an invalid PIM Join/Prune message was received by this device.
99	alaPimRPMappingChange	all	An alaPimRPMappingChange notification signifies a change to the active RP mapping on this device.
100	alaPimInterfaceElection	all	An alaPimInterfaceElection notification signifies that a new DR or DR has been elected on a network.
101	lpsLearnTrap	all	Generated when an LPS port learns a bridged MAC.
102	gvrpVlanLimitReachedEvent	all	Generated when the number of vlans learned dynamically by GVRP has reached a configured limit.
103	alaNetSecPortTrapAnomaly	all	Trap for an anomaly detected on a port.
104	alaNetSecPortTrapQuarantine	all	Trap for an anomalous port quarantine.
105	udldStateChange	all	Generated when the state of the UDLD protocol changes.
106	healthMonIpcTrap	all	This trap is sent when IPC Pools exceed usage.
107	bcmHashCollisionTrap	all	-
108	healthMonCpuShutPortTrap	all	This trap is sent when port is shut down because

No.	Trap Name	Platforms	Description
			of a CPU spike.
109	arpMaxLimitReached	all	This IP Trap is sent when the hardware table has reached the maximum number of entries supported. The OS6400 will not generate new ARP request for new nexthops.
110	ndpMaxLimitReached	all	This IPv6 Trap is sent when the hardware table has reached the maximum number of entries supported. The OS6400 will not generate new ARP request for new nexthops.
111	ripRouteMaxLimitReached	all	This trap is sent when the RIP database reaches the supported maximum number of entries. When the maximum number is reached, RIP discards any new updates.
112	ripngRouteMaxLimitReached	all	This trap is sent when the RIPng database reaches the supported maximum number of entries. When the maximum number is reached, RIPng discards any new updates.
113	aaaHicServerTrap	all	This trap is sent when the HIC server is down.
114	alaErpRingStateChanged	all	This trap is sent when the ERP Ring State has changed from “Idle” to “Protection”.
115	alaErpRingMultipleRpl	all	This trap is sent when multiple RPLs are detected in the Ring.
116	alaErpRingRemoved	all	This trap is sent when the Ring is removed dynamically.
117	e2eGvrpVlanMatch	all	This trap is sent when GVRP recieves a registration for a VLAN that is configured for End-to-End Flow Control.
118	e2eStackTopoChange	all	This trap is sent when the stack topology changes.
119	dot3OamThresholdEvent	all	This trap is sent when a local or remote threshold crossing event is detected. A local threshold crossing event is detected by the local entity, while a remote threshold crossing event is detected by the reception of an Ethernet OAM Event Notification OAMPDU that indicates a threshold event.
120	dot3OamNonThresholdEvent	all	This trap is sent when a local or remote non-threshold crossing event is detected. A local event is detected by the local entity, while a remote event is detected by the reception of an Ethernet OAM Event Notification OAMPDU that indicates a non-threshold crossing event.
121	alaDot3OamThresholdEventClear	all	This trap is sent when is sent when a local or remote threshold crossing event is recovered.
122	alaDot3OamNonThresholdEventClear	all	This trap is sent is sent when a local or remote non-threshold crossing event is recovered.

No.	Trap Name	Platforms	Description
123	ntpMaxAssociation	all	This trap is generated when the maximum number of peer and client associations configured for the switch is exceeded.
124	alaLicenseExpired	9000E	This trap is sent when the value of aluLicenseTimeRemaining becomes 0 (zero) for a demo licensed application. This notification is applicable only for temporary licenses. This trap can be utilized by an NMS to inform user about an application license expiration.
125	vRtrLdpInstanceStateChange	all	This trap is sent when the LDP module changes state either administratively or operationally.
126	vRtrLdpGroupIdMismatch	all	This trap is sent when there is a mismatch of local and remote group IDs.
127	mplsXCup	9000E	This trap is generated when one of the configured cross-connect entries is about to leave the down state and transition into some other state (but not into the “Not Present” state).
128	mplsXCdown	9000E	This trap is sent when one of the configured cross-connect entries is about to enter the down state from some other state (but not from the “Not Present” state).
129	vRtrMplsStateChange	9000E	This trap is sent when the MPLS module changes state.
130	vRtrMplsIfStateChange	9000E	This trap is sent when is generated when the MPLS interface changes state.
131	vRtrMplsLspUp	9000E	This trap is sent when an LSP transitions to the 'inService' state from any other state.
132	vRtrMplsLspDown	9000E	This trap is sent when an LSP transitions out of 'inService' state to any other state.
133	svcStatusChanged	9000E	This trap is sent when there is a change in the administrative or operating status of a service.
134	sapStatusChanged	9000E	This trap is sent when there is a change in the administrative or operating status of an SAP.
135	sdpBindStatusChanged	9000E	This trap is sent when there is a change in the administrative or operating status of an SDP Binding.
136	sdpStatusChanged	9000E	This trap is sent when there is a change in the administrative or operating status of an SDP.
137	sapPortStateChangeProcessed	9000E	This trap is sent when the agent has finished processing an access port state change event, and that the operating status of all the affected SAP's has been updated accordingly.

No.	Trap Name	Platforms	Description
138	sdpBindSdpStateChangeProcessed	9000E	This trap is sent when the agent has finished processing an SDP state change event, and that the operating status of all the affected SDP Bindings has been updated accordingly.
139	unused	-	-
140	unused	-	-
141	unused	-	-
142	ddmTemperatureThresholdViolated	all	This trap is sent when an SFP/ XFP/SFP+ temperature has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/ XFP/SFP+ temperature.
143	ddmVoltageThresholdViolated	all	This trap is sent when SFP/XFP/ SFP+ supply voltage has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ supply voltage.
144	ddmCurrentThresholdViolated	all	This trap is sent when if an SFP/ XFP/SFP+ Tx bias current has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx bias current.
145	ddmTxPowerThresholdViolated	all	This trap is sent when an SFP/ XFP/SFP+ Tx output power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Tx output power.
146	ddmRxPowerThresholdViolated	all	This trap is sent when an SFP/ XFP/SFP+ Rx optical power has crossed any threshold or reverted from previous threshold violation for a port represented by ifIndex. It also provides the current realtime value of SFP/XFP/SFP+ Rx optical power.
147	halHashCollisionTrap	all	-
148	alaLbdStateChangeToShutdown	all	This trap is sent when the port state changes to “shutdown”.
149	alaLbdStateChangeForClearViolationA	all	This trap is sent when the port state changes from “shutdown” due “to clear-violation-all”.
150	alaLbdStateChangeForAutoRecovery	all	This trap is sent when the port state changes from shutdown due to auto-recovery mechanism

No.	Trap Name	Platforms	Description
151	pimBsrElectedBSRLostElection	all	This trap is sent when the current E-BSR loses an election to a new Candidate-BSR.
152	pimBsrCandidateBSRWinElection	all	This trap is sent when a C-BSR wins a BSR Election.
153	alaErpRingPortStatusChanged	all	This trap is sent whenever the ring port status changes.
154	lnkaggPortReserve	all	This trap is sent when given port of the link aggregation goes to reserved state.
155	esmViolationRecoveryTimeout	all	This trap is sent when a user port is re-enabled after an esm viola-tion recovery timeout.
156	alaMvrpVlanLimitReachedEvent	all	This trap is sent when the num-ber of VLANs learned dynami-cally by MVRP reaches the configured limit.
157	alaMvrpE2eVlanConflict	all	This trap is sent when MVRP receives a registration for a VLAN that is configured for End To End Flow Control.
158	alaDhcpSrvLeaseUtilizationThreshold	all	This trap is sent when the lease utilization on a subnet exceeds or falls below the configured threshold value.
159	alaDhcpClientAddressAddTrap	all	This trap is sent when a new IP address is assigned to DHCP Cli-ent interface.
160	alaDhcpClientAddressExpiryTrap	all	This trap is sent when the lease time expires or when the DHCP client is not able to renew/rebind an IP address
161	alaDhcpClientAddressModifyTrap	all	This trap is sent when the DHCP client is unable to obtain the existing IP address and a new IP address is assigned to the DHCP client.
162	alaDyingGaspTrap	all	This trap is sent when a switch has lost all power.
163	alaTestOamTxDoneTrap	all	After a configured time interval, this trap is sent to the NMS from Generator switch when the test duration expires.
164	alaTestOamRxReadyTrap	all	This trap is sent to the NMS once the switch with Analyzer or Loopback Role is ready to receive test traffic. Once this trap is received, the Generator is activated for generating test traffic.
165	alaTestOamTestAbortTrap	all	This trap is sent to the NMS from the switch, if the test is aborted during takeover.
166	Reserved40	-	-
167	Reserved41	-	-

No.	Trap Name	Platforms	Description
168	alaSaaIPIterationCompleteTrap	all	This trap is sent when an IP SAA iteration is completed.
169	alaSaaEthIterationCompleteTrap	all	This trap is sent is sent when a Eth-LB or Eth-DMM SAA iteration is completed.
170	alaSaaMacIterationCompleteTrap		-
171	aaaHicServerChangeTrap	all	This trap is sent when the active HIC server is changed from.to primary.
172	aaaHicServerUpTrap	all	This trap is sent when at least one of the HIC servers comes UP.
173	alaLldpTrustViolation	all	This trap is sent when there is an LLDP Trust Violation, and gives the reason for the violation.
174	alaStackMgrIncompatibleModeTrap	all	-
175	alaEsmDBChange	all	-
176	alaDHLVlanMoveTrap	all	When linkA or linkB goes down or comes up and both ports are are part of some vlan-map, this trap is sent to the Management Entity, with the DHL port information.
177	esmPortViolation	all	This trap is sent when an interface is shut down by a feature due to violation.
178	stpLoopGuardError	all	This trap is sent by a bridge when a port enters the Loop inconsistent state (ERR state).
179	stpLoopGuardRecovery	all	This trap is sent by a bridge when a port leaves the Loop inconsistent state (ERR state).

Unsupported Software Features

CLI commands and Web Management options may be available in the switch software for the following features. These features are not supported in AOS Release 6.4.4.R01:

Feature	Platform	Software Package
OSPF Database Overflow (RFC 1765)	all	base
Authenticated VLANs	OS9000E	base
Legacy VLAN Stacking Mode	all	base
Binding Rules	OS9000E	base
IPX Routing	all	base

Unsupported CLI Commands

The following CLI commands are not supported in AOS Release 6.4.4.R01:

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
BFD	ip bfd-std mode demand
Chassis Mac Server	mac-range local mac-range duplicate-eeprom mac-range allocate-local-only show mac-range status
Chassis Supervision	show fabric
DHCP Relay	ip helper traffic-suppression ip helper dhcp-snooping port traffic-suppression
Ethernet Interfaces	10gig slot [slot] phy-a|phy-b interfaces long interfaces runt interfaces runtsize interfaces flood rate interfaces hybrid preferred-copper interfaces hybrid forced-copper interfaces hybrid forced-fiber
Flow Control	Flow flow wait time interfaces flow
Hot Swap	reload ni [slot] # [no] power ni all
Source IP Management	aaa radius agent preferred ntp src-ip preferred snmp source ip preferred
NTP	no ntp server all
PIM	ip pim cbsr-masklength ip pim static-rp status ip pim rp-candidate ip pim crp-address ip pim crp-expirytime ip pim crp-holdtime ip pim crp-interval ip pim crp-priority ip pim data-timeout ip pim joinprune-interval ip pim source-lifetime ip pim interface mode ip pim interface cbsr-prefernce ip pim interface max-graft-retries

Software Feature	Unsupported CLI Commands
	ip pim interface sr-ttl-threshold show ip pim rp-candidate show ip pim rp-set show ip pim nexthop show ip pim mroute
QoS	qos classify fragments qos flow timeout show policy classify destination interface type show policy classify source interface type
RIP	ip rip redist status ip rip redist ip rip redist metric ip rip redist-filter ip rip redist-filter effect ip rip redist-filter metric ip rip redist-filter route-tag ip rip redist-filter redist-control
System	install show microcode history
VLANs	vlan router mac multiple enable|disable vlan binding mac-port-protocol vlan binding mac-ip vlan binding ip-port show vlan ipmvlan port-binding
VRF	ip service http ip service all
Tunneling L2 Protocols	ethernet-service uni-profile P l2-protocol [STP | GVRP]peer

Unsupported MIBs

The following MIBs are not supported in AOS Release 6.4.4.R01:

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

MIB Name	Unsupported MIB variables
AlcatelIND1Port	esmPortCfgLongEnable esmPortCfgRuntEnable esmPortCfgRuntSize esmPortPauseSlotTime esmPortCfgFLow alcether10GigTable
AlcatelIND1QoS	alaQoSPortPdiTable alaQoSSlotPcamTable alaQoSPortProtocolTable alaQoSSlotProtocolTable alaQoSSlotDscpTable alaQoSRuleReflexive alaQoSAppliedRuleReflexive alaQoSActionSourceRewriteIpAddr alaQoSActionSourceRewriteIpAddrStatus alaQoSActionSourceRewriteIpMask alaQoSActionTable alaQoSActionSourceRewriteNetworkGroup alaQoSActionTable alaQoSActionSourceRewriteNetworkGroupStatus alaQoSActionTable alaQoSActionDestinationRewriteIpAddr alaQoSActionTable alaQoSActionDestinationRewriteIpAddrStatus alaQoSActionTable alaQoSActionDestinationRewriteIpMask alaQoSActionTable alaQoSActionDestinationRewriteNetworkGroup alaQoSActionTable alaQoSActionDestinationRewriteNetworkGroupStatus alaQoSActionTable alaQoSActionLoadBalanceGroup alaQoSActionTable alaQoSActionLoadBalanceGroupStatus alaQoSActionTable alaQoSActionPermanentGatewayIpAddr alaQoSActionTable alaQoSActionPermanentGatewayIpAddrStatus alaQoSActionTable alaQoSActionAlternateGatewayIpAddr alaQoSActionAlternateGatewayIpAddrStatus alaQoSAppliedActionSourceRewriteIpAddr alaQoSAppliedActionSourceRewriteIpAddrStatus alaQoSAppliedActionSourceRewriteIpMask alaQoSAppliedActionSourceRewriteNetworkGroup alaQoSAppliedActionSourceRewriteNetworkGroupStatus alaQoSAppliedActionDestinationRewriteIpAddr alaQoSAppliedActionDestinationRewriteIpAddrStatus alaQoSAppliedActionDestinationRewriteIpMask alaQoSAppliedActionDestinationRewriteNetworkGroup alaQoSAppliedActionDestinationRewriteNetworkGroupStatus alaQoSAppliedActionLoadBalanceGroup alaQoSAppliedActionLoadBalanceGroupStatus alaQoSAppliedActionPermanentGatewayIpAddr alaQoSAppliedActionPermanentGatewayIpAddrStatus alaQoSAppliedActionAlternateGatewayIpAddr alaQoSAppliedActionAlternateGatewayIpAddrStatus alaQoSPortDefaultQueues alaQoSPortAppliedDefaultQueues alaQoSConfigNatTimeout alaQoSConfigAppliedNatTimeout alaQoSConfigReflexiveTimeout alaQoSConfigAppliedReflfexiveTimeout alaQoSConfigFragmentTimeout alaQoSConfigAppliedFragmentTimeout

MIB Name	Unsupported MIB variables
	alaQoSConfigClassifyFragments alaQoSConfigAppliedClassifyFragments
AlcatelIND1Slb	slbFeature slbClusterTable slbServerTableg
AlcatelIND1StackManager	alaStackMgrStatsTable
AlcatelIND1SystemService	systemUpdateStatusTable
AlcatelIND1VlanManager	vlanIpxNet vlanIpxEncap vlanIpxRipSapMode vlanIpxDelayTicks vlanSetMultiRtrMacStatus vlanIpxStatus vlanSetIpxRouterCount
AlcatelIND1WebMgt	alaIND1WebMgtRFSConfigTable alaIND1WebMgtHttpPort alaIND1WebMgtHttpsPort
IEEE_802_1X	dot1xAuthDiagTable dot1xAuthSessionStatsTable dot1xSuppConfigTable dot1xSuppStatsTable
IETF_BGP4	bgpRcvdPathAttrTable bgp bgpPeerTable bgp4PathAttrTabl
IETF_BRIDGE	dot1dTpPortTable dot1dStaticTable
IETF_ENTITY	entLogicalTable entLPMappingTable entAliasMappingTable
IETF_ETHERLIKE	dot3CollTable dot3StatsSQETestErrors dot3StatsInternalMacTransmitErrors dot3StatsCarrierSenseErrors dot3StatsInternalMacReceiveErrors dot3StatsEtherChipSet dot3StatsSymbolErrors dot3ControlInUnknownOpcodes
IETF_IF	ifRcvAddressTable ifTestTable
IETF_IP_FORWARD_MIB	ipForwardTable
IETF_IPMROUTE_STD	ipMrouteScopeNameTable
IETF_MAU (RFC 2668)	rpMauTable rpJackTable broadMauBasicTable ifMauFalseCarriers ifMauTypeList ifMauAutoNegCapability ifMauAutoNegCapAdvertised ifMauAutoNegCapReceived
IETF_OSPF (RFC 1850)	ospfAreaRangeTable

MIB Name	Unsupported MIB variables
IETF_OSPF_TRAP	ospfTrapControl
IETF-PIM	pimRPTable
IETF_P_BRIDGE	dot1dExtBase dot1dPortCapabilitiesTable dot1dPortPriorityTable dot1dUserPriorityRegenTable dot1dTrafficClassTable dot1dPortOutboundAccessPriorityTable dot1dPortGarpTable dot1dPortGmrpTable dot1dTpHCPortTable dot1dTpPortOverflowTable
IETF_Q_BRIDGE (RFC 2674)	dot1qTpGroupTable dot1qForwardAllTable dot1qForwardUnregisteredTable dot1qStaticMulticastTable dot1qPortVlanStatisticsTable dot1qPortVlanHCStatisticsTable dot1qLearningConstraintsTable
IETF_RIPv2	rip2IfConfDomain
IETF_RMON	hostControlTable hostTable hostTimeTable hostTopNControlTable hostTopNTable matrixControlTable matrixSDTable matrixDSTable filterTable channelTable bufferControlTable captureBufferTable
IETF_RS_232 (RFC 1659)	all synchronous and sdlc objects and tables rs232SyncPortTable
IETF_SNMPv2	sysORTable snmpTrap sysORLastChange
IETF_SNMP_ COMMUNITY (RFC 2576)	snmpTargetAddrExtTable
IETF_SNMP_ NOTIFICATION (RFC 2576)	snmpNotifyTable snmpNotifyFilterProfileTable snmpNotifyFilterTable
IETF_SNMP_PROXY (RFC 2573)	snmpProxyTable
IETF_SNMP_TARGET (RFC 2573)	snmpTargetAddrTable snmpTargetParamsTable snmpTargetSpinLock
IETF_SNMP_USER_BASED_SM (RFC 2574)	UsmUser
IETF_SNMP_VIEW_BASED_ACM (RFC 2575)	vasmMIBViews

Open Problem Reports and Feature Exceptions in Release 6.4.4.R01

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

SNMP

PR	Description	Workaround
157020	PoE connect and disconnect traps are received only on the initial disconnect. Subsequent disconnects do not generate a trap.	There is no known workaround at this time.

LAYER 2

Ethernet OAM

PR	Description	Workaround
156081	efm-oam l1-ping does not work with dynamic link aggregation ports.	Configure static link aggregation ports.

Dual-Home Link Aggregation

PR	Description	Workaround
157874	If port 1/1 is configured as part of a Dual-Home Link Aggregate (Active-Active), either as a physical port or part of a link aggregate, the default VLAN cannot be changed on any other port in the switch.	Do not configure port 1/1 as part of a Dual-Home Link Aggregate (Active-Active).
152732	After removing VLANs using the 'no vlan' command the VLANs are not removed from the Dual-home Link vlan-map.	Manually remove the VLANs from the DHL configuration

LLDP

PR	Description	Workaround
153023	In some circumstances an LLDP port may be moved to the "violation" state before the default violation timer interval of (3 * LLDP transmit interval).	There is no known workaround at this time.
153696	If a port is operationally down or LLDP trust-agent is disabled the state of the port displays as TRUSTED".	There is no known workaround at this time.

Source Learning

PR	Description	Workaround
152080	When Port-based Ingress Source Filtering is enabled, the VLAN-based ISF will be flushed and re-programmed along with Port based ISF config, this allows some packets to flow temporarily during re-programming.	There is no known workaround at this time.
152082	When VLAN-based Ingress Source Filtering is enabled, the VLAN-based ISF will be flushed and re-programmed, this allows some packets to flow temporarily during re-programming.	There is no known workaround at this time.

LAYER 3

BGP

PR	Description	Workaround
156500	Unable to ping IPv6 neighbor after entering 'no ip bgp bestpath med missing-as-worst' command and resetting the ports.	There is no known workaround at this time.

Security

Access Guardian

PR	Description	Workaround
157990	LPS configuration is removed and a boot.cfg.err generated when upgrading from 6.4.3 to 6.4.4 for the following command due to CLI parameter change: `->  port-security <slot/port> enable`	Re-enter the command using the new CLI parameter: `->  port-security <slot/port> admin-status enable`
157480	HIC host operational status shows as 'bridging' instead of 'HIC' when client is in HOLD mode.	There is no known workaround at this time. This is a display issue only, HIC functionality works as expected.
157739	The "show 802.1x users" command displays the status of a failed supplicant as "authenticating", even though the client is in a blocked state.	Use the "show mac-address-table" command to verify the MAC address for the supplicant client is in a filtering state.

Port Mirroring/Monitoring

PR	Description	Workaround
151905	On an OmniSwitch 9000E when port monitoring is configured on an egress port only the unmodified ingress BOOTP/DHCP unicast routed packets will be monitored if DHCP relay is not configured.	User port mirroring

OPEN PROBLEM REPORTS FROM PREVIOUS RELEASES

PR	Description	Workaround
95308	Temporary traffic loops could happen under the following scenarios: 1. Reloading of a non root bridge. This happens when the bridge is going down and is due to the sequential bringing down of NIs during a reload process .It is purely temporary in nature and stops when all the NIs eventually get powered off. 2. NI power down When an NI power down command is executed for an NI and if that NI has the Root port port and other NIs have Alternate ports, it is possible to see some traffic looping back from the newly elected Root port. The traffic loop back is temporary and will stop once the NI gets powered off. 3. New Root bridge selection Temporary loops could occur during the process of electing a new Root bridge, if this election process is triggered by the assignment a worse priority for the existing root bridge or a root bridge failure. This happens due to the inconsistent spanning tree topology during the convergence and stops entirely once the network converges	For items 1 and 2 above there is no work around presently. For item 3 the following work around could be applied: 1. Tune the max age (and or max hops in the case of MSTP) parameter to a lower value that is optimal for the network. This will reduce the convergence time and thereby the duration of temporary loops. 2. To select a new root bridge, consider assigning better priority for that bridge instead of assigning worse priority for the existing root bridge.
111029	The 'show mac-address-table count' command may not display the correct number of learned MAC entries for link aggregation ports after an STP event.	There is no known workaround at this time. This is a display issue only.
113928	After a MAC movement due to a new mobility rule match the entry may still be displayed with the previous information.	There is no known workaround at this time. This is a display issue only.
138770	In a stacked environment on a takeover where the NI is reset, the polling frame from the switch does not reach the supplicant.	There is no known workaround at this time.
143071	Sometimes an AVLAN MAC address doesn't get removed from the CLI display when using the 'show mac-address-table' command.	This is a display issue only, the MAC address is correctly removed from the system. Use the 'show avlan user' command to correctly display the AVLAN MAC addresses.
145589	On an OS6850 auto-negotiation configuration needs to be replicated on both fiber and copper mediums for combo ports.	Use the following commands to duplicate the auto-negotiation configuration: -> interfaces <slot/port> hybrid fiber autoneg {enable | disable} -> interfaces <slot/port> hybrid copper autoneg {enable | disable}
145667	When configuring VPLS with 4K SAPs, SDP status may remain down after a switch reload.	There is no known workaround at this time.

Hot Swap / Redundancy

Feature Exceptions

CMM and Power Redundancy Feature Exceptions for OmniSwitch

- Manual invocation of failover (by user command or Primary pull) should only be done during times when traffic loads are minimal.
- Hot standby redundancy or failover to a secondary CMM without significant loss of traffic is only supported if the secondary is fully flash synchronized with the contents of the primary's flash.
- Hot standby redundancy or failover to a secondary module without significant loss of traffic is only supported if all the remaining units in the stack are fully flash synchronized with the contents of the primary's flash.
- Failover/Redundancy is not supported when the primary and secondary CMMs are not synchronized (i.e., unsaved configs, different images etc.). In this case, upon failover, all the NIs will reset and might go to "down" state, and to recover, need to power down the switch and power it back up.
- Primary and Redundant power supplies must be of the same type. For example, having a primary 510W power supply with a redundant 360W power supply is not supported.

Hot Swap Feature Exceptions for OmniSwitch 9000E

- Hot swap of NIs needs to be preceded by the removal of all cables connected to the NI.
- Hot swap of unlike modules is not supported.
- The **reload ni** command is not supported. Please use **no power ni**/**power ni** as an alternative.
- All insertions of NI modules cannot be followed by another hot swap activity until the OK2 LED on the inserted NI blinks green.

Hot Swap Feature Exceptions for OmniSwitch 6400/6850(E)/6855

- When removing modules from the stack (powering off the module and/or pulling out its stacking cables), the loop back stacking cable must be present at all times to guarantee redundancy. If a module is removed from the stack, rearrange the stacking cables to establish the loopback before attempting to remove a second unit.
- When inserting a new module in the stack, the loop back has to be broken. Full redundancy is not guaranteed until the loop back is restored.
- Hot swap of unlike chassis is not supported.

Hot Swap Time Limitations for OmniSwitch

- 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.

- All insertions of stack elements must be done one at a time and the inserted element must be fully integrated and operational as part of the stack before inserting another element.

Technical Support

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

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

Email: esd.support@alcatel-lucent.com

Web: service.esd.alcatel-lucent.com

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

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

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

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

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

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